AI-Absorbing Innovation*

Adair Morse
University of California - Berkeley
& NBER

Amrita Nain
University of Iowa

March 2026

Abstract

An unstudied avenue by which AI impacts employment and economic growth is through the assimilation of artificial intelligence (AI) into non-technology innovation. Using patent citations to identify instances where non-IT patents incorporate AI knowledge, we document an increasing dominant U.S. market share in AI-absorbing innovation and an excess market valuation. Our primary goal is to study the implications of AI-absorbing innovation for employment and its firm mechanisms. Using Jorda (2005) projections, Sun and Abraham (2021) estimators, and instrumented methods using laboratory (federal and university) AI, we find that AI-absorbing innovation increases employment. Isolating the mechanism, firms engaging in AI-absorbing innovation display faster growth in output as well as the profits and needed capital associated with the revenue growth. We find that the positive effects of AI-absorbing innovation are not primarily through automation. Rather output growth is driven by both AI-absorbing product innovation and AI-absorbing process innovation, reflecting channels identified in Babina et al (2024), but with a very different AI role.

Keywords: Employment effects from Artificial Intelligence (AI), patent valuation, innovation dissemination, automation and labor, innovation gap, AI and firm growth

* Author contacts: adair@berkeley.edu and amrita-nain@uiowa.edu. Nain thanks the financial support of University of Iowa, Tippie College of Business for research grant support. We thank Renping Li and Gerald Hoberg for helpful comments.

I. Introduction

Artificial intelligence (AI) is widely expected to generate fundamental shifts in the economy, triggering an active debate on AI's effects on employment, largely through channels of human capital substitution, effort augmentation, and productivity enhancement.[1] A separate strand of literature focuses on innovation itself, speaking to the technology industry and its growth in driving macroeconomic and market valuation growth, particularly in the United States.[2] Yet, with a few notable exceptions, these literatures have largely developed in parallel, with little attention to how advances in AI enhance innovation in other sectors.[3] In this paper, we study the extent to which innovation outside the information technology (IT) space assimilates AI and the impact of this AI-absorbing innovation on employment and output growth. Non-IT patents account for 80% of patenting in the United States, and thus until we understand the role of AI on non-IT innovation, we cannot speak to the aggregate economic consequences to employment, growth, and productivity.

We begin with an example to fix ideas. Our example is a 2021 microsurgery device patent (US 10,895,742). The device uses optical imaging to capture the surgical field and applies AI-driven real-time image processing to isolate and display magnified views of the area of surgical interest. This patent builds on earlier advances in biometrics and facial-recognition technologies that improved image identification precision (2012 patent US 8,315,440 and 2014 patent US 8,824,779). By absorbing AI technologies, the device can dynamically adjust to changing visual conditions during surgery, improving precision and adaptability across procedures and enabling remote surgical guidance and telemedical support.

The point to the example is this: AI absorption need not be classified as an AI technology. Furthermore, the economic impact of such AI absorption may arise through innovation itself rather than through labor substitution or the hiring of AI specialists. While prior work shows that AI hiring can stimulate product innovation (Babina et al., 2024) and that automation may alter labor demand (Acemoglu et al., 2022; Hampole et al., 2025), AI may also enhance firms' inventive capabilities by expanding the scope and functionality of domain-specific innovations. In this sense,

---

[1] Acemoglu, Autor, Hazell, and Restrepo, 2022; Babina, Fedyk, He, and Hodson, 2024; Hampole, Papanikolaou, Schmidt, and Seegmiller, 2025; Brynjolfsson, Chandar, and Chen, 2025.
[2] Gazzani and Natoli, 2024
[3] See for example, Babina, et al, 2024 and Fedyk, Mihet, Rishabh, and Gomes, 2024. We discuss our contribution relative to these works.

AI absorption can act as a catalyst for further innovation rather than merely as a tool that augments or replaces labor. This is the subject of our work herein.

There is little systematic evidence on the degree to which AI knowledge is absorbed into innovation outside the technology sector. In light of this, our study puts forth four contributions. First, we document how prevalent AI-absorbing innovation is in non-technology sectors of the U.S. economy and show that such innovation generates excess economic value. Second, we estimate a plausibly causal relationship between AI-absorbing innovation and higher employment, with the usual caveats on causal interpretation. Third, we provide evidence that the value created by AI-absorbing innovation arises primarily through output expansion and associated profitability growth rather than through efficiency or productivity improvements. Finally, we show that AI-absorbing innovation drives top-line growth through both product and process innovation.

We identify AI-absorbing innovation using patent citations. This approach follows a large literature that interprets patent citations as proxies for knowledge spillovers and technological knowledge flows (e.g., Jaffe, 1986; Jaffe, Trajtenberg, and Henderson, 1993; Hall, Jaffe, and Trajtenberg, 2001). Our analysis uses the universe of patents granted by the United States Patent and Trademark Office (USPTO) between 1979 and 2024. Using International Patent Classification (IPC) codes, we identify AI patents and measure AI-absorbing innovation as instances in which non-IT patents cite these AI patents. We document a steady increase in AI-citing patents over our sample period, with U.S. firms especially dominant in this type of innovation. In the last two years of our sample period, about 10% of non-IT patents issued to U.S. firms are AI-citing. Moreover, the United States holds a 50-percentage-point larger share of AI-absorbing patents than Europe. This gap is moreover growing, which is especially noteworthy because the US-Europe gap in total innovation and in AI innovation has remained relatively flat over this period. Using market valuation techniques of Kogan, Papanikolaou, Seru, and Stoffman (2017), and controlling for firm, and patent class times year fixed effects, we find non-IT patents that are AI-absorbing in citing an AI patent have 40 basis point higher, significant patent value.[4]

With this descriptive set of first contributions in hand, we turn to the agenda to bring AI-absorbing innovation into the AI-employment literature, which is carrying a seemingly large voice

[4] Interestingly, within the subsample of IT patents, citing an AI patent has no effect on patent market value. One interpretation of these findings is that AI is a general-purpose technology whose value emerges when it is combined with domain-specific knowledge from other sectors, rather than through incremental improvements within the originating technology sector (Bresnahan and Trajtenberg, 1995).

in welfare consequences. The concern is that firms may reduce labor demand by replacing workers with AI applications or AI-powered machines. Consistent with this possibility, prior studies using granular task-level data find that AI-exposed tasks experience lower labor demand (Hampole et al., 2025). However, the impact of AI on occupation- or firm-level employment appears more ambiguous. Because AI and automation can increase productivity and firm growth, they may also raise the demand for labor (Acemoglu and Restrepo, 2019; Babina et al., 2024; Hampole et al., 2025). In light of this debate, we examine whether AI-absorbing innovation increases or decreases employment.

Extending our patent-level data to firm-year prevalence of AI-citing patents, we estimate local projection regression (LPR) models of Jorda (2005). We find that a one standard deviation increase in AI-absorbing innovation results in 0.23% higher employment per year over a five-year horizon. This estimation faces the standard endogeneity concerns in that the AI-absorbing innovation may reflect firm selection on the availability of profitable innovation opportunities and on their stock of knowledge capital, which enables them to recognize and successfully exploit the potential value of AI-related technologies.

We implement three strategies to address these concerns. First, we restrict the sample to innovative firms, comparing firms that engage in AI-absorbing innovation with innovative firms that do not. The positive employment effects remain qualitatively similar. Second, we estimate an interaction-weighted (IW) treatment estimator of Sun and Abraham (2021) adapted to include an anticipation window of the patenting process (a la Callaway and Sant'Anna, 2021). In this modern two-way fixed effects (TWFE) specification, we improve on LPR assumptions using staggered treatment and a more precise definition of the dynamic control comparison. Consistent with the baseline results, we continue to find a significant positive effect of AI-absorbing innovation on employment.

Even under this more stringent design, our estimates may still reflect firms' knowledge capital: their awareness of AI-related opportunities and their ability to translate them into value through AI-absorbing innovation. Such a possibility is especially true when the treatment is a human decision. Hence, our third strategy to address endogeneity is to implement an instrumental variable strategy that exploits localized knowledge spillovers, abstracting from firm decisions altogether. Specifically, we instrument firms' AI-absorbing innovation with the local stock of AI patents produced by universities and federal laboratories in the same commuting zone building off

Jaffe, Trajtenberg, and Henderson (1993) and Helper, Makan, and Shoag (2025). Conditional on industry and region controls, the local supply of AI knowledge from academic and government research institutions is plausibly unrelated to firm-specific shocks affecting outcomes, satisfying the exclusion restriction. Although most inventors are located in the same commuting zone as the assignee firm, roughly 20% of patents involve inventors located elsewhere. This creates within–commuting-zone variation in the instrument across firms, allowing us to include commuting zone times year fixed effects. In our IV estimations, we continue to find evidence of a significant positive impact of AI-absorbing innovation on employment.

Our third contribution involves documenting the mechanisms behind the effect of AI-absorbing innovation on labor. One possibility is that AI-absorbing innovation raises productivity that enhances labor demand. Alternatively, AI-absorbing innovation may drive growth through the creation of new products and services, increasing output and employment even without productivity gains. To tease out these explanations, we study firm output, COGS/Sales, gross profit, total factor productivity (TFP) and capital stock. Across our different methodologies, we find robust evidence that AI-absorbing innovation leads to higher growth in output, profit, and capital stock. Critically for pinning down the mechanism, our evidence rejects any significant treatment effects between AI-absorbing innovation and efficiency or productivity.

Overall, our results suggest that the value-enhancing effects of AI-absorbing innovation are driven primarily by revenue growth. In traditional R&D theory, product innovation drives revenue expansion, while process innovation targets cost reductions (Cohen and Klepper, 1996; Klepper, 1996). However, recent work by Baslandze, Liu, Sojli, and Tham (2025) argues that process innovation can also stimulate growth by unlocking technologies and efficiencies that enable future product innovation. Motivated by these arguments, we examine whether the pro-growth effects of AI-absorbing innovation are due to product innovation, process innovation, or both.

Following Bena, Ortiz-Molina, and Simintzi (2022) and Bena and Simintzi (2025), we classify patents as product or process innovations using the text of patent claims. We then construct firm-level measures of AI-absorbing product and process innovation and re-estimate our baseline regressions. We find that AI-absorbing product innovation is associated with growth in output, employment, profits, and capital, but not with TFP. These findings are consistent with Babina et al. (2024), who show that hiring AI talent increases product innovation. Importantly, our measure of AI-absorbing innovation is distinct from AI hiring: using their data, we find only a modest correlation (about 0.09) between the two measures, and our results remain unchanged when controlling for the number of AI employees.

We also find a positive role for process innovation. AI-absorbing process innovation is associated with higher growth in output, profits, and employment, although the effects on output and profits emerge with a delay, appearing only in the third through fifth year following the innovation. This pattern is supportive of the Baslandze et al (2025) argument that the effect of process innovation on growth is through *future* product innovation. Overall, our findings relating to process innovation are different from Babina et al (2024) who find no link between AI hiring and process innovation. We note that our results do not contradict their evidence but complement it as our measure of AI-absorbing innovation is distinct from AI-hiring.

Our final set of tests further confirm that the value generated by AI-absorbing innovation is not primarily driven by labor-replacing automation. While textual analysis of patents reveals that 35% of AI-absorbing patents relate to automation, we find that non-automation AI-absorbing innovation strongly predicts growth in employment, output, profit, and capital. In contrast, AI-absorbing "smart automation" shows weaker and less consistent effects. Overall, these findings suggest that the economic value of AI absorption arises primarily through non-automation channels.

Our paper contributes to the growing literature studying the economic consequences of artificial intelligence by shifting attention from the production of AI technologies to their assimilation into innovation outside the IT sector. Existing work largely studies either the labor market effects of AI adoption (e.g., Acemoglu et al., 2022; Babina et al., 2024; Hampole et al., 2025) or the role of AI innovation in driving technological progress and growth. By contrast, we focus on how non-IT firms incorporate AI knowledge into their own inventive activity and show that this "AI-absorbing innovation" is both prevalent and economically valuable.

Our paper also contributes to the large literature studying the effects of automation on labor markets. Much of this literature emphasizes how technological change alters the demand for labor by automating tasks previously performed by workers. Early work on skill-biased technological change shows that new technologies often increase the relative demand for skilled labor (e.g., Acemoglu 2002; Autor, Katz, and Krueger 1998). More recent research develops a task-based framework in which automation substitutes for labor in some tasks while creating new tasks in others (Autor, Levy, and Murnane 2003; Acemoglu and Restrepo 2019). Empirical work using data on robots, AI adoption, and online vacancies studies how these technologies affect employment, wages, and task composition (Acemoglu and Restrepo 2020; Acemoglu et al. 2022; Babina et al. 2024). While this literature has generated important insights into the labor-displacing and labor-augmenting effects of automation, it primarily focuses on the adoption of automation technologies within firms or occupations**.** We highlight a complementary channel through which AI may affect labor demand: the assimilation of AI knowledge into the innovation activities of firms outside the IT sector**.**

More broadly, the paper contributes to the literature on absorptive capacity and knowledge spillovers. Prior research emphasizes that firms' ability to recognize, assimilate, and exploit external knowledge is central to innovation (Cohen and Levinthal, 1989; Jaffe, 1986; Zahra and George, 2002; Audretsch and Belitski, 2020). We extend this literature by providing new evidence on how firms combine distinct knowledge domains—specifically AI and domain-specific technologies—to generate new innovative capabilities. By documenting how AI knowledge diffuses across technological fields through patent citations, and that U.S. firms lead Europe in this process, we highlight a new mechanism through which absorptive capacity may contribute to differences in innovation and productivity across economies.

The rest of the paper is organized as follows. Section II discusses the data and describes our measure of AI-absorbing innovation. Section III presents results relating to the employment effects of AI-absorbing innovation. Section IV studies the mechanisms through which AI-absorbing innovation creates value. Section V concludes.

II. AI Absorbing Patents

We begin by describing the construction of the measure of AI absorbing patents and motivating its importance in the economy, highlighting its role in the US-Europe innovation gap. In the final part of this section, we assess the market value of AI absorbing patents using the Kogan et al (2017) estimator.

II.a Patent Data

We download 8.4 million utility patents granted by the U.S. Patent and Trademark Office (USPTO) between 1979 and 2024 from USPTO PatentsView.[5] The dashed line in Figure 1.A depicts the growth of patenting over grant year. We identify the location of patent assignees and plot the number of patents for U.S. based assignees (blue line) and Europe-based assignees (red dotted line) in the same figure,[6] focusing on these two geographies for a practical reason of motivation. Figure 1.C depicts the US-Europe innovation gap defined as the share of patents granted to U.S. firms minus the share of patents granted to European firms. The innovation gap (i.e., U.S. advantage) ranges from 0.3 to 0.4 over the sample period, with a declining pattern since the height in the mid-1990s.

In Figure 1.B, we depict the total number of artificial intelligence (AI) patents granted by USPTO, as well as the subset granted to U.S. and European firms over our sample period.[7] Figure 1.D shows that the US-Europe innovation gap in AI innovation, which is noticeably stronger in AI innovation than in total innovation and which has stayed remarkably steady at 0.5.

II.b Measurement and Importance of AI Absorbing Patents

We follow standard practice in capturing knowledge assimilation through patent citations. Since our objective is to examine how firms outside the technology sector assimilate AI in their innovative capability, we measure AI absorbing as the extent to, if any, which a current patent cites a prior patent that is classified as AI. We are interested in the absorbing of AI knowledge across the economy (i.e., in industry, agriculture, non-tech services, etc.) rather than the absorbing of AI knowledge into information technology (IT) advancements themselves. Thus, we intentionally exclude IT patents as the current or focal patents from our sample.[8]

Figure 2.A plots 6.7 million non-IT patents granted by USPTO over our sample period, while Figure 2.C depicts the difference in share of non-IT patents of US and Europe. We highlight two facts from this graph. First, non-IT patents account for almost 80% of patents granted by

---

[5] Utility patents account for 90% of USPTO patents and the term is used to distinguish from design patents. Note that in all graphs, the final year or two is more volatile from the flow of reporting and may not reflect the final counts in the period.
[6] If a patent has an assignee in more than one country, the patent is allocated to both countries
[7] AI patents are defined as patents with International Patent Classification (IPC) codes G06N3, G06N5, G06N7, G06N10, G06F15, and G06K9. See appendix Table A2 for a description of these IPC codes.
[8] IT patents are patents in the following IPC classes - G06F, G06Q, G06N, G06K, G06T, H04L, H04W, H04M. Note that all AI classes listed in Footnote 7 are nested within IT classes. See appendix Table A2 for a description of these IPC codes.

USPTO. Second, while US firms maintain an innovative advantage in the non-IT patent classes throughout the sample period, the gap has shrunk since the late 1990s from about 0.4 to 0.2. This is a remarkable departure from the story told in Figures 1.C and 1.D.

Figure 2.B depicts 300,865 patents that absorb AI granted by USPTO in the non-IT patent classes. In post-2022, 10% of non-IT patents are AI-citing. Figure 2.D plots the US-Europe innovation gap in patents that absorb AI. In stark contrast to the downward trend in Figure 2.C, Figure 2.D shows that US firms have increased their advantage relative to Europe in innovation that absorbs AI from about 0.4 in the 1990s to about 0.5 at the end of our sample period.

Already we believe this pattern to be an important new stylized fact: the innovation gap of the US over Europe is increasing in the absorption of AI into the economy.[9] To the extent that AI absorbing innovation contributes to firm growth and productivity, this advantage in AI absorption implies consequences for the future of these two economies.

II.c Value of AI Absorbing Patents

To study any differential value implied by the assimilation of AI knowledge relative to other patents, we use the patent market value (ξ) technique of Kogan et al (2017). We create two patent-level variables. The first is *PctAIcite*, calculated as the number of backward citations to AI patents divided by the total number of backward citations. The second is I_{AIcite}, an indicator variable that takes the value one if at least one backward citation is made to an AI patent. We restrict our sample of patents to those granted to U.S. publicly traded firms and require the availability of the IPC classification of both the focal patent and at least one cited patent.

In Panel A of Table 2, we summarize PctAIcite, I_{AIcite}, $\ln(\xi)$ and other patent characteristics for 2.5 million patents (including IT patents) issued to publicly traded firms in the US between 1979 and 2024. All variables are described in Appendix A. The mean value of I_{AIcite} indicates that about 20% of the patents cite at least one AI patent. Many of these are likely to be cases where the

[9] AI-absorbing patents are identified using citations made to U.S. granted patents only. While USPTO PatentsView provides data on citations to foreign patents, the identification numbers of the foreign patents are not standardized and cannot be matched easily to IPC classes. The exclusion of foreign citations raises concerns that our AI-citing measure suffers from a home bias in knowledge assimilation. That is, when applying for USPTO patents, European firms may be citing AI patents issued by the European Patenting Office (EPO) instead of by USPTO. This would lead to a downward bias in AI-absorbing patents for European firms. To address this concern, we download USPTO patents from the PATSTAT Global database instead of from PatentsView. The advantage of the PATSTAT data is that the IDs of foreign cited patents are standardized. We find that the advantage of US firms in AI-absorbing patents persists even when all foreign citations are included.

focal patent is an IT patent (including focal AI patents) that cites an AI patent. Since our interest is in assimilation of AI into innovation conducted outside the IT sector, we henceforth focus on Panel B excluding focal IT patents from the sample. In non-IT patents, the mean I_{AIcite} is 0.067; just under 7% of non-IT patents cite at least one AI patent. The mean value of PctAIcite appears low at 1.2%, lower than the indicator because PctAIcite captures the yes/no on citing any AI plus the percentage of patents cited being AI. If we look within the subset of non-IT patents with I_{AIcite}=1 (i.e. at least one AI citation), the mean value of PctAIcite is 18% (not tabulated). We note that the rest of the patent characteristics such as $\ln(\xi)$, forward and backward citations, firm size, volatility are similar across Panel A and B.

We estimate the relation between AI absorption and the market value of patent p granted to firm i as follows:[10]

$$ln(\xi_p) = \alpha + \beta\ PctAIcite_p + \gamma_1 X_i + \gamma_2\ ln\, Backward_p + \mu_i + \rho_{kt} + \epsilon_p, \quad (1)$$

where $ln\, Backward$ is the log of the number patents cited by the patent, included because assimilation of any existing knowledge, including non-AI knowledge, can be valuable. X is a vector of firm-level control variables including log market capitalization (patenting may increase with size) as of the day prior to patent grant and log idiosyncratic volatility (mechanically correlated with ξ and reflective of high-growth firms having more volatile returns and higher patenting value). Finally, we include technology class times grant year fixed effects, ρ_{kt}, to force the comparison of AI assimilation to within technology-year compariables and further include μ_i, firm-fixed effects, for make the comparison within the firm's norm.[11]

As shown in Table 2, we find that AI absorbing patents are not more valuable than other patents in general (column (1)) or for IT patents (column (3)). However, reported in column (2), non-IT AI absorbing patents command higher economic value than other patents, significant at the 1% level.[12] The inference from the β coefficient in column (2) is that an AI-citing patent commands a 40 basis points higher patent value ξ if the patent cites AI at the mean value (0.18) of PctAIcite when positive AI citing happens relative to a patent that does not cite any AI.[13]

---

[10] In Table B1 of Appendix Bm we use the indicator variable I_{AIcite} instead of PctAIcite.
[11] Results are qualitatively similar if we use firm times year fixed effects and technology class fixed effects instead (see appendix Table B2)
[12] In the appendix Table B1, we find qualitatively similar results if we use I_{AIcite} as the measure of AI absorption.
[13] All else equal, the change in $ln(\xi)$ if PctAIcite changes from 0 to the average of 0.18 is 0.18*β, which equals 0.004 using β = 0.022 from column 3 of Table 2. Since the dependent variable is the natural log of ξ, this represents a 0.4% increase in ξ.

Our results show that the market values transformative capability as defined by Zahra and George (2002) – i.e. the ability of firms to innovate by synthesizing two unconnected sets of knowledge. Combined Table 2 and Figure 2.D show that AI^2 outside the technology space is value-enhancing and that U.S. firms are increasingly engaging in this value-enhancing innovative activity.

III. AI^2 and Employment

The rest of the paper seeks to understand why AI-absorbing patents are highly valued by the market. We begin by exploring the effect of AI^2 on employment. While AI can automate some tasks and reduce labor demand, it also permits workers to increase productivity by reallocating effort toward tasks that are difficult to automate (Hampole et al, 2025; Acemoglu et al, 2022). With these competing forces in mind, our objective is to understand whether AI^2 benefits a firm through reduction in labor (i.e. lower employment growth) or by presenting growth opportunities leading to more labor demand. Along the way, we analyze how AI^2 implementation does and does not overlap the known patent channels – automation and process innovation.

III.a Data at the Firm Level

To study firm-level outcomes, we construct a firm measure of AI absorbing innovation (henceforth denoted AI^2) as the percentage of patents in a firm-year that cite at least one AI patent. (This is equivalent to the firm-year average of the patent-level indicator variable I_{AIcite}.) A second measure, which we use for robustness, is an indicator variable called $I_{AI\text{^}2}$ equal to one for firm-years with at least one patent that cites an AI patent, and zero otherwise. Since we are interested in understanding the effects of assimilating AI knowledge outside of the technology sector, we exclude firms in technology sectors.[14] Our measure differs from the existing AI-employment nexus literature by focusing on innovation explicitly, by focusing on non-IT sectors, and by focusing on citing of AI rather than AI patenting directly. Other novel AI measures include Babina et al (2024), who measure human capital AI via the hiring of AI graduates and Hampole et al, (2025) and Brynjolfsson et al (2018), who extract task-level AI exposures via job description data.

[14] Information technology firms are defined as firms with the following three-digit SIC codes - 357, 481, 482, 484, 489, and 737.

Table 3 Panel A summarizes AI^2 and I_{AI^2} for the full sample while Panel B summarizes the variables for an innovative-only subsample. We define innovative firms as those with at least one patent in the period *t* through *t-5*. This second sample, our preferred, allows us to make comparisons in among comparably innovating firms.

In the full sample, 2.6% of the stock of patents in a firm-year cite at least one AI patent, with an average of I_{AI^2} of 0.067, implying that just under 7% of firm-years have at least one focal patent that cites an AI patent. These percentages are low because majority of firm-years in our sample have no patents. In the innovative sample, the mean of AI^2 is 0.075 – almost 8% of the stock of patents in a firm-year cite at least one AI patent – with nearly 20% of firm-years having at least one AI absorbing focal patent.

III.b Methodology Overview

Our estimation faces common endogeneity concerns of corporate actions. AI^2 actions may reflect a selection based on characteristics of the firm and its setting and on intentional firm effort. Specifically, firms doing frontier AI^2 are a selection on both opportunities from innovation as well as a firm's knowledge capital in knowing and being able to unlock positive expected value with AI^2. Our identification approach in such a setting takes a series of steps. Our first two methods work to refine a dynamic comparison of other firms with equally innovating trajectories in the spirit of the advancements to difference-in-difference under the two-way fixed effects (TWFE) literature (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; Goodman-Bacon, 2021; Borusyak, Jaravel, and Spiess, 2024). However, one might argue that this comparison alone is insufficient in that AI^2 innovating firm-years may still reflect a selection on a firm's dynamic knowledge capital as to the utility of AI^2. To address this possible knowledge capital endogeneity, we instrument AI^2 through a local knowledge capital spillover channel. We follow Jaffe et al (1993) and Helper, Makan, and Shoag (2025) to capture local spillovers from universities and federal labs.

III.c Jordà (2005) Local Projection Analysis

Our first methodology follows Kogan et al. (2017), Kelley, Papanikolaou, Seru, and Taddy (2021) and Baslandze et al. (2024) in estimating the effect of AI^2 on employment using the Jordà (2005) Local Projection Regressions (Jordà LPR) methodology.

III.a.1 Local Projection Methodology

Following this literature, we estimate firm *i*'s employment response to AI^2 at various horizons starting forward from time *t* according to the model:

$$\ln(Emp_{i,t+\tau}) - \ln(Emp_{it}) = \alpha_0 + \beta_\tau AI{}^\wedge 2_{i,t} + \gamma_1 X'_{i,t} + \gamma_2 \ln(Emp_{i,t-1}) + \delta_{j,t} + \epsilon_{i,t}. \quad (2)$$

Emp is the number of employees; our dependent variable is the change in the natural log of employment for a firm over horizons τ of $\tau = 1, \ldots, 5$ years. Our main dependent variable, AI^2, is the fraction of the firm's patents filed in that year that are AI-citing. Identification requires isolating the AI^2 effect relative to an appropriate, dynamic control group. We include lagged log initial employment $Emp_{i,t-1}$, thus estimating log changes off initial employment. We level dynamically to firm size and risk by including the log value of capital stock and log idiosyncratic volatility in vector X_{it}. We also include the log of one plus the number of patents a firm is granted in the year as a measure of innovative capacity. Finally, we isolate the dynamic control group with the industry-year by including SIC3 times year fixed effects δ_{jt}, double clustering standard errors by firm and year.

III.a.2 Local Projection Results

Table 4 presents the LPR results. In panel A, we include all firms in the estimation, within the SIC3 times year fixed effect. We first pause to consider the control variables. We find the usual relationship that employment growth declines in percentage change as firms grow larger. The coefficient on lagged log employment is negative across the columns of the projection. By contrast, the employment elasticity to patenting is robustly positive but small. At a one-year horizon (column 1), a 10% higher patenting implies a 0.06% larger employment growth.

Our main interest is to consider the coefficient on AI^2. Reading the projection effect across from left to right on the table, we find a positive AI^2 effect on employment, with statistically significant coefficients ranging from 0.029 for one-year projection to 0.89 for cumulative five-year projection. Interpreting these economic magnitudes relative to summary statistics in Table 3, we find that a standard deviation higher AI^2 (0.132 higher AI^2) results in a 0.38% higher employment for one year, or 0.23% annual higher employment if simple averaged over five years.

Arguably, however, the effect of AI^2 should be compared only to other innovating firms. Thus, in panel B, we restrict the sample to only innovating firms in the firm-year observation, where we define innovating as having a positive patent application in the year. We find that this comparable group precision matters, reducing the magnitude of the impact of estimated AI^2 on employment and increasing the coefficient on log number of patents. Using the coefficients on AI^2 and the same standard deviation applied for panel A magnitude (0.132), we find that a standard deviation higher AI^2 implies a 0.22% higher annual employment if simple averaged over five years.

III.b Sun and Abraham (2021) Dynamic Treatment Analysis

The LPR results presented above may be subject to endogeneity concerns raised by the recent advancements in the two-way fixed effects (TWFE) models by Wooldridge (2021), Goodman-Bacon (2021), Callaway and Sant'Anna (2021), Sun and Abraham (2021), Borusyak, Jaravel, and Spiess (2024), and others. Two of the concerns are as follows. First, when applied to panel data with staggered treatment and heterogenous treatment effects, standard LPR estimates can suffer from bias because these models compare later treated units to previously treated units. This can cause the coefficients β_τ in equation (1) to be contaminated by the effects of already units.

Second, AI^2 does not happen immediately, but may reflect an evolving company strategy that ultimately results in a patent grant. Such a setting would violate, for example, the Abbring and van den Berg (2003) and Sianesi (2004) treatment anticipation assumption.

III.b.1 Dynamic Treatment Methodology

To address these concerns, we estimate dynamic treatment effects using the Sun and Abraham (2021) Interaction-Weighted (IW) estimator with anticipation horizon assumptions from Callaway and Sant'Anna (2021). Unlike the Jordà LPR method that pools all treated cohorts into a single average, the IW estimator explicitly isolates each treatment cohort. By estimating cohort-specific effects and then aggregating them using cohort sample shares as weights, this method ensures that "already-treated" units are not used as controls. We identify when a firm is first treated

– i.e. when a firm first engages in AI absorbing innovation – as the filing year of the firm's first patent that cites an AI patent,[15] thus defining:

- *Treatment Cohort (G_{ig}):* $G_{ig} \in \{g_1, g_2, \ldots, g_n\}$, an indicator for a firm being in the cohort first treated at time g, and
- *Control Group* (C_i)*:* $C_i \in \{0, 1\}$, the never treated.

Since the AI^2 strategy and innovation likely begin a few years before the patent application, we allow for an anticipation window, a "strategy and innovation phase." Following Callaway and Sant'Anna (2021), such a period would parallel trends. Thus, to satisfy limited anticipation, conditional on covariates and the horizon, we assume:

- *Anticipation Window* (*$\delta \geq 0$*): the number of anticipation periods before the treatment date where effects begin to manifest, and
- *Reference Period* (*$g - \delta - 1$*): The baseline period omitted from the regression to ensure a "clean" untreated comparison.

We can now introduce the outcome language:

- *Potential Outcome* ($Y_{it}(g)$), the potential outcome for firm *i* at time *t* if it were treated at time g,
- *Untreated Potential Outcome* ($Y_{it}(0)$), the counterfactual outcome for firm *i* at time *t* if it were never treated, and
- *Observed Outcome* (Y_{it}): $Y_{it} = \sum_{g \in G} G_{ig} Y_{it}(g) + C_i Y_{it}(0)$, the observed outcome in the data for firm *i* at time *t*.

Our dynamic treatment effects assumptions are threefold:

1. Conditional Parallel Trends (CPT)**:**
$$E\left[Y_{it}(0) - Y_{g-\delta-1}(0) \middle| X, G_g = 1\right] = E\left[Y_t(0) - Y_{g-\delta-1}(0) \middle| X, C = 1\right] \tag{3}$$
2. Limited Anticipation with Window δ:
$$Y_{it} = Y_{it}(0) \;\text{ if }\; G_{ig} = 1 \text{ and } t < g - \delta \tag{4}$$
3. Treatment Effect (ATT):
$$\tau(g, t, X) = E\left[Y_t(g) - Y_t(0) \middle| X, G_g = 1\right] \neq 0 \quad \text{for } t \geq g - \delta \tag{5}$$

The Sun and Abraham method involves first estimating cohort average treatment effect on the treated at each period. The Sun and Abraham (2021) Interaction-Weighted (IW) estimator is:

$$Y_{it} = \alpha_i + \lambda_t + \sum_{g \in \mathcal{G}} \sum_{\ell \neq -\delta-1} \theta_{g\ell}\left(G_{ig} \cdot \mathbb{1}\{t - g = \ell\}\right) + \varepsilon_{it}, \tag{6}$$

[15] In keeping with our norm of studying the non-IT sectors of the economy, we require that the firm not be a software or communications firm (SIC codes 357, 481, 482, 484, 489, 737).

where $\mathbb{1}\{t - G_{ig} = \ell\}$ is an indicator for firm i at time t being ℓ periods away (the *relative time*) from its initial treatment time g. α_i and λ_t are firm and time fixed effects respectively. The coefficient estimator $\hat{\theta}_{gl}$ from equation (2) is a difference-in-difference estimator of the cohort average treatment effect on the treated. The Sun-Abraham IW estimator is then calculated for each relative period ℓ by taking a weighted average of $\hat{\theta}_{gl}$ using sample shares of each cohort in the relevant period ℓ as weights.

III.b.2 Dynamic Treatment Results

Our initial task is to assert a window for anticipation and then to check graphically for the validity of no anticipation. The literature on the gestation of patents (Mansfield, Schwartz, and Wagner, 1981; Mansfield, 1988; Griliches, 1990; Cohen, Nelson, and Walsh, 2000; Heeley, Matusik, and Jain, 2007) finds that the patent development period ranges from 5 months to 4 years, on average, depending on the time period and industries studied. We choose a middle range of two years.

Focusing on log employment as our initial outcome Y_{it}, we plot the IW estimates for 10 leads and lags as Figure 3. The figure strongly suggests anticipation in the two-year window. Although if one wants to consider the impact of the t-10 observation, perhaps there is some trend in effect, but this is not statistically supported nor would it have been evident if we had chosen a different lead-lag range other than 10. Hence, our inference is that the limited anticipation of Callaway and Sant'Anna (2021) is not rejected in our setup.

We are now able to interpret the dynamic treatment model that uses the staggered design, has permanent treatment, and discards the shaded anticipation window. These specification differences between the LPR and IW methodologies are material. Yet, despite the entirely different design and the added stringency of our IW method, we find an effect that is slightly muted but very similar to the economic magnitude inference from Table 4, panel B. This result can be seen in Figure 3. In the staggered post period, the treatment of AI^2 results in a 0.17% increase in employment. Furthermore, this positive employment effect is permanent, in the sense the effect is one-time and persistent over the horizon studied.

III.c Knowledge Capital Endogeneity & IV Identification through Laboratory AI

Despite the plausibility of the assumptions we made in the dynamic treatment analysis, one might argue that what we are identifying is not the AI^2 patenting itself. We might be picking up an intentional firm strategy toward a growth agenda, where the impact of the strategy causes the employment effect. In particular, AI^2 innovating firm-years may reflect a selection on a firm's dynamic knowledge capital as to the utility of AI^2.

III.c.1 Laboratories AI: IV Assumptions

In this section, we deploy the linear propagation-instrumental variable (LP-IV) methodology of Ramey (2016). Our IV strategy seeks to step away from knowledge capital decisions inside the firm. Specifically, we instrument a firm's AI^2 with the local stock of AI patents (not AI absorbing patents) generated by universities and federal laboratories in the same commuting zone (CZ). This strategy is motivated by Jaffe, Trajtenberg, Henderson (1993) and Helper et al (2025), who document that knowledge spillovers are highly localized geographically, with firms disproportionately citing patents produced nearby.

Universities and federal laboratories produce frontier scientific knowledge. We argue that this knowledge production is exogenous to individual firm innovation decisions *other than* possibly the activities in the technology sector. A higher local stock of AI patents from these public-sector and university institutions therefore increases a firm's exposure to AI knowledge through localized spillovers, allowing us to use this measure as an instrument for AI^2, under the empirical question of whether the first stage relevance holds. Our exclusion restriction is that conditional on region and industry controls, the historical accumulation of AI patents in nearby academic and government research institutions is plausibly unrelated to contemporaneous firm-specific shocks affecting outcomes except through the firm's use of AI knowledge.

We calculate the local AI knowledge stock (local AI stock) for patent p granted in year t as the total number of AI patents granted to universities and federal laboratories located in the same commuting zone as the inventor of patent p over the five years from t-5 to t. If a patent has multiple inventors, we use the highest value of local AI stock among its inventors. We then collapse the patent level measure local AI stock to the firm-level by taking the mean value of this across all patents in a firm-year. We name this firm-level variable *Laboratory AI*. While the majority of inventors are in the same location as the assignee (i.e. the firm), for about 20% of patents, the inventor location differs. This allows variation in the *Laboratory AI* for firms within the same commuting zone. We estimate this Linear Projection-IV (LP-IV), defined by same LPR setup of Table 4 including a vector of firm-level controls of log value of capital stock, log idiosyncratic volatility, ln(1+$Patents_{i,t-1}$), log ($Emp_{i,t-1}$) and $\delta_{j,t}$ -- SIC3 times year fixed effects. We additionally add $\psi_{c,t}$, commuting zone times year fixed effects.

III.c.2 Laboratory AI: IV Results

Results of the LP-IV analysis of employment growth are presented in Table 5. Panels A and B present results for the full sample and innovative sample respectively. The first-stage Kleibergen-Paap (KP) F-statistics across all column horizons and in both panels are greater than 200. The local stock of patents granted to universities and federal labs have relevance as an instrument for AI^2.

In the second stage, the coefficient on AI^2 is positive and statistically significant across all short-term horizons in both panels, with some loss of power at the 5-year horizon in the innovative sample. The magnitude of the estimates are a ten-fold larger effect than those of Table 4. Overall, using plausibly exogenous variation in AI^2, we continue to find evidence of employment growth.

IV. AI^2 and Firm Mechanisms

Our results thus far indicate that AI^2 is valued by the market and that it leads to employment growth. To understand why, we turn to the standard economic arguments of the role of AI in the workplace. AI^2 could be labor augmenting through increases in productivity (Hampole et al, 2025). When AI automates some tasks allowing an existing level of workers to switch effort toward revenue growth-oriented activities. This line of thinking of improvements in productivity can drive economic growth reflects the long-standing understanding of capital-labor productivity models (e.g., Solow, 1957; Romer, 1990; Aghion and Howitt, 1992).[16]

However, as importantly documented by Babina et al (2021), AI^2 can also lead to growth through the creation of new products and services. Thus, growth in output and employment might be observed without growth in capital-use productivity. We can explore which of these mechanism

[16] See literature on skill-biased technological change previously cited in Section IIIc

is at work driving the AI^2 results by exploring LPR estimations for output, COGS/Sales, profits, TFP, and capital.

Important in such analysis is our distinguishing our results and story from the Babina et al (2024) and Hampole et al (2025) channels of firm increasing their AI investment through hiring of AI employees. Yet, the Babina et al measure of AI employee investment has only a limited correlation of 0.10 with our measure of AI^2. This not only suggests that AI^2 is a distinct phenomenon and deserves a separate investigation, but it also suggests that the product innovation insights of Babina et al (2024) may be at play through other mechanisms that these authors had envisioned.

IV.a Output & Productivity Results

We turn to the estimation of mechanisms drivers of the AI^2-employment results. We repeat our analysis flow of first estimating LPR results, then plotting our results from an extended Sun and Abraham (2021) IW estimator with Callaway and Sant'Anna (2021) assumptions of limited anticipation, and finally estimating an IV specification.

Table 6 follows the specifications of Kogan et al. (2017), Kelley et al. (2021) and Baslandze et al. (2024) in estimating the effect of AI^2 on firm performance outcomes in the Jorda form of the log of the horizon length change in the output. These firm performance dependent variables follow the literature, including gross output, cost of goods sold (COGS) scaled by sales, gross profit, total factor productivity (TFP), and capital stock. On the right-hand side, we always include the log lag of the dependent variable plus log capital stock, and log idiosyncratic volatility. Finally, we include SIC3 times year fixed effects. Variable definitions are included as Appendix A and the summary statistics for these variables appear in Table 3. As in our prior Jorda-LPR estimations, we take two approaches for controlling for innovativeness. First, we present the results for the full sample, controlling for the log of one plus the number of patents, (panel A) and then for the sample restricted to innovative firms (panel B). We focus on discussion on Panel B, more conservative in comparability assumptions and in results.

As reported in Table 6, we find that AI^2 projects growth forward positively in gross output and profits. We find weaker positive effects for capital stock growth. We find no robust effect on efficiency of output (COGS/Sales) or productivity. Before considering the economic magnitude and meaning of these results, we investigate the robustness of these first findings in the Sun and

Abraham IW estimation figures – Figure 4, panels A to E – and then in the IV. In Figure 4, panel D, we first rule out a TFP effect. The weak TFP results in Table 6 are indeed not robust to the IW specification. If anything, the treatment of AI^2 results in a propulsion of a slightly negative TFP. The remaining figures are consistent with Table 6. Output (Figure 4A), profits (Figure 4C), and capital (Figure 4E) exhibit positive firm growth from AI^2 treatment, and AI^2 has no treatment effect on efficiency (COGS/sales).

Finally, we turn to the LP-IV specifications, using the Laboratory AI instrument for AI^2. Table 7 presents the results. We continue to find strong evidence for an impact of AI^2 on profits and capital growth, consistently across the treatment projection horizon. Our output results in panels A and B are slightly less consistently statistically significant. Yet, interpreting panel B, the more stringent specification in only innovative firms, we find evidence that AI^2 causes revenue growth at least in the shorter horizon forward.

Taken together, the results paint a picture of the mechanism driving our employment finding. Firms who are innovating in a way that absorbs AI in breakthroughs that have industrial uses (as opposed to technological advancements) has increases in future employments that stems from firm growth in their top lines – revenues and assets, -- not in their efficient or productive use of these assets. The picture here is that firms are growing new product lines or unlocking customers because of AI improvements.

IV.b Overlap with Process Patenting & Automation & The Nature of AI^2

Our results suggest that the value-enhancing effects of AI-absorbing innovation operate primarily through an output growth channel leading to higher revenues, profits, and need for capital stock. In this section, we take up whether the nature of these results stems from concepts of process patenting and automation and how our results contribute to these literatures.

IV.b.1 Motivation of Process Patenting & Automation Drivers

Process innovation has traditionally been viewed as a driver of lower production costs and labor-saving efficiency, both of which appear to be less relevant mechanisms in our findings.[17]

[17] Bena and Simintzi (2025) find that US firms reduced process innovation after the 1999 U.S.-China bilateral agreement that increased access to cheaper labor. Bena, Ortiz-Molina, and Simintzi (2022) increase process innovation when state-level legal changes increase labor dismissal costs. Earlier work by Cohen and Klepper (1996) and Klepper (1996) views process innovation as reducing average production costs.

Yet, recent evidence suggests that process innovation can be a revenue-driver. In particular, Baslandze et al. (2025) show that process innovation can lead to sustained firm growth by improving product quality and by enabling the development of new products over time. This perspective is relevant in light of the findings of Babina et al. (2024) that firms' investments in AI hiring increase product innovation but have little effect on process innovation. Yet our mechanism of innovation via AI-absorbing actions may represent a channel distinct from - and potentially complementary to - the AI-hiring channel emphasized in Babina et al. (2024).

Likewise, a natural nexus exists between automation and labor-saving efficiency, again less relevant in our findings.[18] Yet, a channel from automation to revenues is easily found in the literature. For example, Autor et al (2003) and Graetz and Michaels (2018) show that although computers and robotics substitute for labor in routine tasks, they complement workers in non-routine, problem-solving tasks. This increase of skilled worker demand is reflected many times over in the literature[19] and can lead to output growth.

IV.b.2 Process Patenting & Automation Data & Statistics

Motivated by these arguments, we construct data on process patenting and automation. To identify process patents, we follow the method of Bena et al. (2022) and Bena and Simintzi (2025) and identify process patents based on a textual analysis of patent claims (see Appendix C for a description of process patent classification). A patent with more than 50% process claims is classified as a process patent while the rest are classified as product patents. Using this approach, about 35% of AI-citing patents outside the IT space are process patents. To identify automation patents, we use natural language processing to classify each patent in our sample as an automation patent based on the textual description of the patent. A detailed description of the classification of a patent as an automation patent is presented in Appendix C.

In panel A of Table 8, we document the raw statistics: of all non-IT patents, whereas 6.7% are AI citing, automation and process patenting represent 13.9% and 28.5% respectively, indicating that automation and process patenting are much broader concepts than AI^2.

[18] Autor et al (2003); Acemoglu and Restrepo (2019)

[19] Autor and Dorn, 2013; Berman, Bound, and Griliches, 1994; Berman, Bound, and Machin, 1998; Acemoglu, 1998; Autor, Katz, and Krueger, 1998

Panels B and C look at the overlap of automation and process patents explaining the indicator variable of an AI absorbing patent. Rather than report raw correlations (largely driven by the zero no patenting firms), we instead estimate the explanatory power (partial R-Squared) of automation and process patents in explaining the variation of AI absorbing patents (in panel B) and of AI^2 (at the firm level, in panel C), allowing us to control for technology or industry times year fixed effects. In the high-dimensional fixed effects specifications in column (2) of Panel B, the partial R-squared of 0.007 indicates the automation indicator uniquely explains only 0.7% of the remaining 'within-firm' and 'within-technology-year' variation in AI-citing patents. Similarly, specifications in columns (4) of Panel B, the partial R-squared of 0.001 indicates the process indicator uniquely explains only 0.1% of the remaining variation in AI-citing patents. At the firm level in Panel C, column (2) and column (4), the partial R-squared indicates that after controlling for the total number of patents and including industry times year fixed effects, automation and process patents explain about 14% and 4% of the variation in AI^2, respectively.

IV.b.3 Process versus Product Patenting AI^2 Results

The analysis in Table 8 indicates that AI^2 is not simply another proxy for previously studied notions of process innovation or automation. Yet, there is overlap between AI^2 and process innovation. Looking within non-IT patents only, we find that about 35% of AI-citing patents are process patents. This points toward possible complementarity whereby AI^2 generates value through process improvements. In this section, we examine whether the pro-growth effects of AI-absorbing innovation act through product innovation, process innovation, or both. We run LPR again regressions using the following two measures of AI^2 as the main regressors - *Product AI^2* defined as the fraction of a firm's patents that are AI-citing product patents and *Process AI^2* defined the fraction of a firm's patents that are AI-citing process patents. We control for the firm's overall focus on process innovation by include the variable *Process* which is the fraction of a firm's patents that are process patents of any type.

Results are presented in Table 9 for the full sample. Results for the innovative subsample are qualitatively the same but not tabulated for brevity. In panel A, we see that Product AI^2 and Process AI^2 both project forward employment growth positively. In the first period, the result is statistically weak for Process AI^2, but it strengthens in the longer time periods. Panels B and D show that Product AI^2 projects forward output and profit growth positively and this strong positive effect is immediately visible. Process AI^2 also projects positive output and profit growth but the positive effect of process innovation is delayed, showing up significantly only in periods t+3 to t+5. This delay is consistent with the Baslandze et al (2025) argument that the impact of process innovation on growth works through expanding possibilities for *future* product innovation.

The positive impact of Process AI^2 also provides a useful complement to the findings of Babina et al., who show that AI hiring is not associated with increased process innovation. While their results suggest that AI specialists primarily contribute to product-oriented innovation, our evidence allows for the interpretation that process AI innovations are implemented by domain-specific employees—such as engineers, production specialists, or technicians—who integrate AI tools into operational workflows rather than by newly hired AI specialists themselves. That is, the growth benefits of AI may arise not only from hiring AI talent, but also from the ability of existing technical personnel to embed AI technologies into firm-specific production processes.

IV.b.4 Automation Results

Our finding that AI-absorbing innovation generates top line growth through product AI^2 and through process AI^2, suggests that AI-absorption into the broad economy is not just about the automation of tasks. Yet smart automation that incorporates AI^2 may be at work, allowing us to connect recent literature on AI with the seminal literature on routine automation.[20]

We check this by analyzing whether AI-citing patents are automation oriented through a text-analysis of the patent descriptions. We find that about 35% of AI-citing patents outside the technology space meet the definition of automation patents. Although this ratio is similar to the percentage of AI-citing process patents, there is little overlap between the two types of AI-citing patents. Only about 1.5% of AI-citing patents outside the IT space meet the definition of both process patents and automation patents.

We examine the extent to which the pro-growth effects of AI-absorbing innovation act through smart automation, defined as AI-citing automation patents innovation, process innovation, or both. We run LPR again regressions using the following two measures of AI^2 as the main regressors: *Automation AI^2* is our measure of smart automation, defined as the fraction of a firm's

[20] For recent work on the effect of AI, see Acemoglu, Autor, Hazell, and Restrepo, 2022; Hampole, Papanikolaou, Schmidt, and Seegmiller, 2025; Brynjolfsson, Chandar, and Chen, 2025. For the long-standing literature on routine automation, see Autor, Levy and Murnane (2003), Acemoglu and Restrepo (2019), Acemoglu and Restrepo (2020) and references therein.

patents that are AI-citing automation patents and *Non-automation AI^2* is defined the fraction of a firm's patents that are AI-citing non-automation patents. We control for the variable *Automation*, which is calculated as the fraction of a firm's patents that are automation patents of any type. This control variable captures any routine, non-AI based automation patents the firm develops.

Results are presented in Table 10. We first note that Non-Automation AI^2 has a strong positive effect on growth in employment, output, profit, and capital. This further corroborates the message in Table 9 that the effects of AI^2 are not just about automation. The positive effects of Automation AI^2, i.e. of smart automation, are somewhat sparse. We can see a positive effect on employment in the t+3 and t+5 horizons. The effect on output is only weakly positive over these two horizons. The effect of Automation AI^2 on profit is weakly positive over several horizons. We see no effect of smart automation on growth in capital, TFP or COGS/Sales. Interestingly, however, the variable Automation, that includes routine automation, has a strong positive effect on growth in employment, profit over several horizons and strong positive effect on capital over the t+1 and t+2 horizons.

Overall, our results in section IVb suggest that AI^2 primarily generates value through non-automation avenues such as process and product innovation. The pro-growth effects of automation are visible more strongly in routine automation than in smart automation.

V. Conclusion

Artificial intelligence is widely viewed as a transformative technology with potentially large implications for innovation, productivity, and employment. Existing research largely focuses on the development of AI technologies themselves or on the adoption of AI within firms through hiring and task automation. In this paper, we study a complementary channel: the assimilation of AI knowledge into innovation conducted outside the information technology sector. We refer to this as AI-absorbing innovation and measure it using citations by non-IT patents to prior AI patents.

We document that U.S. firms appear to hold an increasing advantage relative to European firms in this dimension of innovation. Using the patent value framework of Kogan et al. (2017), we show that non-IT patents that absorb AI knowledge are more valuable than comparable patents, consistent with the idea that combining AI with domain-specific technologies generates economically meaningful innovations.

To understand the higher value of AI-absorbing patents, we examine the implications of AI-absorbing innovation for firm outcomes. Across several empirical approaches—including local projection regressions, dynamic DID estimators, and an instrumental-variable strategy based on localized AI knowledge spillovers—we find that firms engaging in AI-absorbing innovation experience higher employment growth. We also find that AI-absorbing innovation predicts higher growth in output, profits, and capital stock, but does not generate robust improvements in costs or total factor productivity. These findings suggest that the economic value of AI-absorbing innovation arises primarily through revenue expansion rather than cost-saving efficiency gains.

Finally, we show that both product-oriented and process-oriented AI-absorbing innovations contribute to these outcomes. Product AI innovations generate immediate growth effects, while process AI innovations produce delayed but economically meaningful gains. Overall, our results highlight the importance of understanding how general-purpose technologies such as artificial intelligence diffuse across technological domains and shape firm growth and employment.

References

Abbring, Jaap H., and Gerard J. van den Berg. 2003. “The Nonparametric Identification of Treatment Effects in Duration Models.” Econometrica 71 (5): 1491–517.

Acemoglu, Daron. 1998. “Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality*.” The Quarterly Journal of Economics 113 (4): 1055–89.

Acemoglu, Daron. 2002. “Technical Change, Inequality, and the Labor Market.” Journal of Economic Literature 40 (1): 7–72.

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo. 2022. “Artificial Intelligence and Jobs: Evidence from Online Vacancies.” Journal of Labor Economics 40 (S1): S293–340.

Acemoglu, Daron, and Pascual Restrepo. 2019. “Automation and New Tasks: How Technology Displaces and Reinstates Labor.” Journal of Economic Perspectives 33 (2): 3–30.

Acemoglu, Daron, and Pascual Restrepo. 2020. “Robots and Jobs: Evidence from US Labor Markets.” Journal of Political Economy 128 (6): 2188–244.

Aghion, Philippe, and Peter Howitt. 1992. “A Model of Growth Through Creative Destruction.” Econometrica 60 (2): 323–51.

Audretsch, David B., and Maksim Belitski. 2020. “The Role of R&D and Knowledge Spillovers in Innovation and Productivity.” European Economic Review 123 (April): 103391.

Autor, David H., and David Dorn. 2013. “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market.” American Economic Review 103 (5): 1553–97.

Autor, David H., Lawrence F. Katz, and Alan B. Krueger. 1998. “Computing Inequality: Have Computers Changed the Labor Market?*.” The Quarterly Journal of Economics 113 (4): 1169–213.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. “The Skill Content of Recent Technological Change: An Empirical Exploration.” The Quarterly Journal of Economics 118 (4): 1279–333.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2024. “Artificial Intelligence, Firm Growth, and Product Innovation.” Journal of Financial Economics 151 (January): 103745.

Baslandze, Salomé, Leo Liu, Elvira Sojli, and Wing Wah Tham. 2025. “Foundational Processes and Growth.” FRB Atlanta Working Paper, FRB Atlanta Working Paper, February 19, 2025–01.

Bena, Jan, Hernán Ortiz-Molina, and Elena Simintzi. 2022. “Shielding Firm Value: Employment Protection and Process Innovation.” Journal of Financial Economics 146 (2): 637–64.

Bena, Jan, and Elena Simintzi. 2025. “Machines Could Not Compete with Chinese Labor: Evidence from US Firms’ Innovation.” Review of Finance 29 (6): 1619–61.

Berman, Eli, John Bound, and Zvi Griliches. 1994. “Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufacturers.” The Quarterly Journal of Economics 109 (2): 367–97.

Berman, Eli, John Bound, and Stephen Machin. 1998. "Implications of Skill-Biased Technological Change: International Evidence*." The Quarterly Journal of Economics 113 (4): 1245–79.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2024. "Revisiting Event-Study Designs: Robust and Efficient Estimation." The Review of Economic Studies 91 (6): 3253–85.

Bresnahan, Timothy F., and M. Trajtenberg. 1995. "General Purpose Technologies 'Engines of Growth'?" Journal of Econometrics 65 (1): 83–108.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. "Generative AI at Work*." The Quarterly Journal of Economics 140 (2): 889–942.

Callaway, Brantly, and Pedro H. C. Sant'Anna. 2021. "Difference-in-Differences with Multiple Time Periods." Journal of Econometrics, Themed Issue: Treatment Effect 1, vol. 225 (2): 200–230.

Cohen, Wesley M., and Steven Klepper. 1996. "Firm Size and the Nature of Innovation within Industries: The Case of Process and Product R&D." The Review of Economics and Statistics 78 (2): 232–43.

Cohen, Wesley M., and Daniel A. Levinthal. 1990. "Absorptive Capacity: A New Perspective on Learning and Innovation." Administrative Science Quarterly 35 (1): 128–52.

Cohen, Wesley M., Richard Nelson, and John P. Walsh. 2000. Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not). Working Paper No. 7552. NBER Working Papers. National Bureau of Economic Research, Inc.

Fedyk, Anastassia, Roxana Mihet, Kumar Rishabh, and Orlando Gomes. 2024. "<div> Data Innovation Complementarity and Firm Growth." SSRN Scholarly Paper No. 4559921. Social Science Research Network, September 1.

Gazzani, Andrea Giovanni, and Filippo Natoli. 2024. "The Macroeconomic Effects of AI-Based Innovation." SSRN Scholarly Paper No. 4938014. Social Science Research Network, December 5.

Goodman-Bacon, Andrew. 2021. "Difference-in-Differences with Variation in Treatment Timing." Journal of Econometrics, Themed Issue: Treatment Effect 1, vol. 225 (2): 254–77.

Graetz, Georg, and Guy Michaels. 2018. "Robots at Work." The Review of Economics and Statistics 100 (5): 753–68.

Griliches, Zvi. 1990. "Patent Statistics as Economic Indicators: A Survey." Journal of Economic Literature 28 (4): 1661–707.

Hall, Bronwyn H., Adam Jaffe, and Manuel Trajtenberg. 2005. "Market Value and Patent Citations." The RAND Journal of Economics 36 (1): 16–38.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller. 2025. "Artificial Intelligence and the Labor Market." Working Paper No. 33509. Working Paper Series. National Bureau of Economic Research, February.

Heeley, Michael B., Sharon F. Matusik, and Neelam Jain. 2007. “Innovation, Appropriability, and the Underpricing of Initial Public Offerings.” The Academy of Management Journal 50 (1): 209–25.

Helper, Susan, Resem Makan, and Daniel Shoag. 2025. Federally-Funded Science as Engines of Regional Innovation and Prosperity: The Case of the National Labs. Working Paper.

Jaffe, Adam B. 1986. “Technological Opportunity and Spillovers of R & D: Evidence from Firms’ Patents, Profits, and Market Value.” The American Economic Review 76 (5): 984–1001.

Jaffe, Adam B., Manuel Trajtenberg, and Rebecca Henderson. 1993. “Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations.” The Quarterly Journal of Economics 108 (3): 577–98.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. 2021. “Measuring Technological Innovation over the Long Run.” American Economic Review: Insights 3 (3): 303–20.

Klepper, Steven. 1996. “Entry, Exit, Growth, and Innovation over the Product Life Cycle.” The American Economic Review 86 (3): 562–83.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. “Technological Innovation, Resource Allocation, and Growth*.” The Quarterly Journal of Economics 132 (2): 665–712.

Mansfield, Edwin. 1988. “The Speed and Cost of Industrial Innovation in Japan and the United States: External vs. Internal Technology.” Management Science 34 (10): 1157–68.

Mansfield, Edwin, Mark Schwartz, and Samuel Wagner. 1981. “Imitation Costs and Patents: An Empirical Study.” The Economic Journal 91 (364): 907–18.

Ramey, Valerie. 2016. “Macroeconomic Shocks and Their Propagation.” In Handbook of Macroeconomics, vol. 2. Elsevier.

Romer, Paul M. 1990. “Endogenous Technological Change.” Journal of Political Economy 98 (5): S71–102.

Sianesi, Barbara. 2004. “An Evaluation of the Swedish System of Active Labor Market Programs in the 1990s.” The Review of Economics and Statistics 86 (1): 133–55.

Solow, Robert M. 1957. “Technical Change and the Aggregate Production Function.” The Review of Economics and Statistics 39 (3): 312–20.

Sun, Liyang, and Sarah Abraham. 2021. “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects.” Journal of Econometrics, Themed Issue: Treatment Effect 1, vol. 225 (2): 175–99.

Wooldridge, Jeffrey. 2021. Two-Way Fixed Effects, the Two-Way Mundlak Regression, and Difference-in-Differences Estimators.

Zahra, Shaker A., and Gerard George. 2002. “Absorptive Capacity: A Review, Reconceptualization, and Extension.” The Academy of Management Review 27 (2): 185–203.

Figure 1: Total Innovation and AI Innovation

Figure 1.A. depicts the total number USPTO patents over the period 1979 to 2024 and the number of patents granted to U.S firms and European firms respectively. Figure 1.C. depicts the difference in the share of total patents held by U.S. and European companies, i.e., the difference between the blue (solid line) share of all patents and the red (dashed) line share of all patents. Figure 1.B. depicts the number of Artificial Intelligence (AI) patents granted by USPTO, and the number of AI patents granted to U.S. and European firms respectively. Patents with IPC class G06N3, G06N5, G06N7, G06N10, G06F15, and G06K9 are classified as AI patents. Figure 1.D. depicts the difference in the share of total AI patents held by U.S. and European companies.

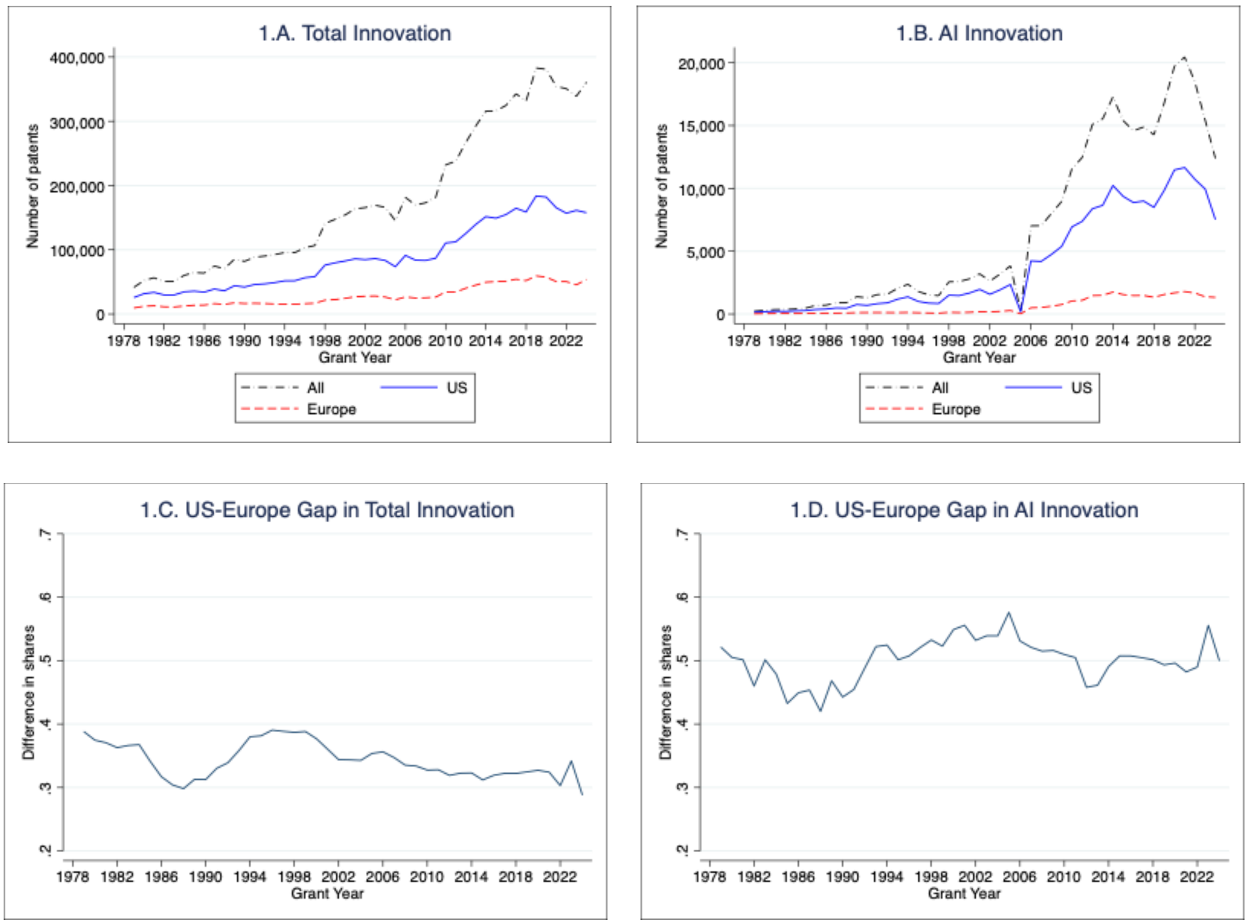

Figure 2: AI-Absorbing Innovation in the Non-IT Sector

Figure 2.A. depicts the total number non-IT patents granted by USPTO over the period 1979 to 2024 and the number of non-IT patents granted to U.S firms and European firms respectively. Non-IT patents are patents *not* in the following information technology IPC classes: G06F, G06Q, G06N, G06K, G06T, H04L, H04W, H04M. Figure 2.C. depicts the difference in the share of not-IT patents held by U.S. and European companies, i.e., the difference between the blue (solid line) share of all patents and the red (dashed) line share of all patents. Figure 2.B. depicts AI-absorbing innovation, defined as non-IT patents that cite AI patents. Patents with IPC class G06N3, G06N5, G06N7, G06N10, G06F15, and G06K9 are classified as AI patents. Figure 1.D. depicts the difference in the share of AI-absorbing innovation by U.S. and European companies.

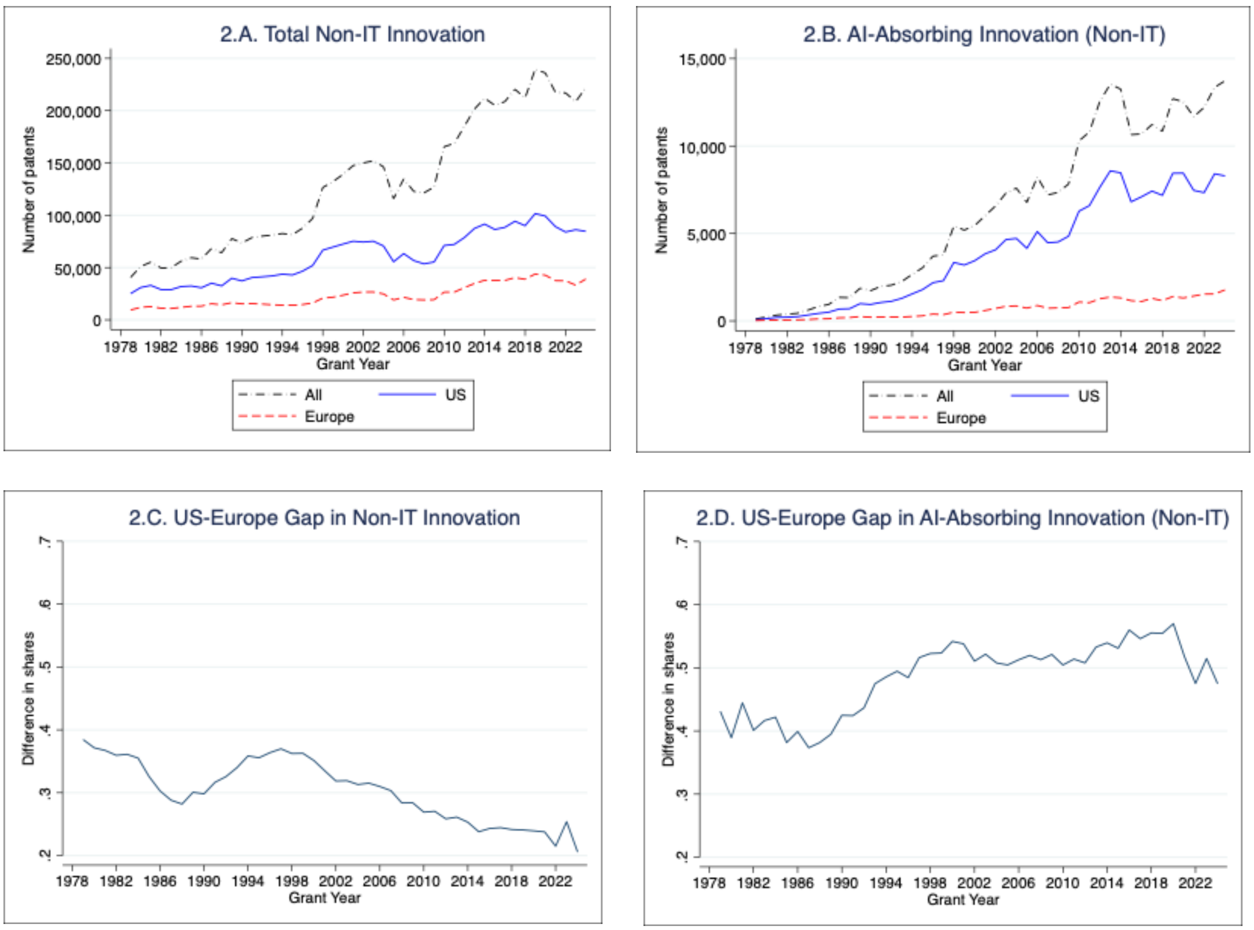

Figure 3
Two-Way Fixed Effects Analysis: Employment

This figure depicts the Sun and Abraham (2021) Interaction-Weighted (IW) estimator for ten periods before and ten periods after first treatment. First treatment is the filing year of the first AI-absorbing (AI^2) patent by a firm. Using G to denote the year when a unit first becomes treated, a cohort average treatment effect on the treated, $\hat{\theta}_{g,l}$, is first estimated using the following regression

$$Y_{it} = \alpha_i + \lambda_t + \sum_{g} \sum_{\ell \neq -3} \theta_{g\ell} \big(G_{ig} \cdot \mathbb{1}\{t - g = \ell \,\big) + \varepsilon_{it}$$

Y is log employment and $G_{ig} = \mathbb{1}(t - g = \ell\,)$ is an indicator for unit i being ℓ periods away from initial treatment at calendar time t. α_i are firm-fixed effects and λ_t are year-fixed effects. The figures show, for each relative period ℓ, the weighted average of $\hat{\theta}_{gl}$ using sample shares of each cohort in the relevant period ℓ as weights. The shaded region represents the period of anticipatory behavior during which AI^2 innovation is already underway.

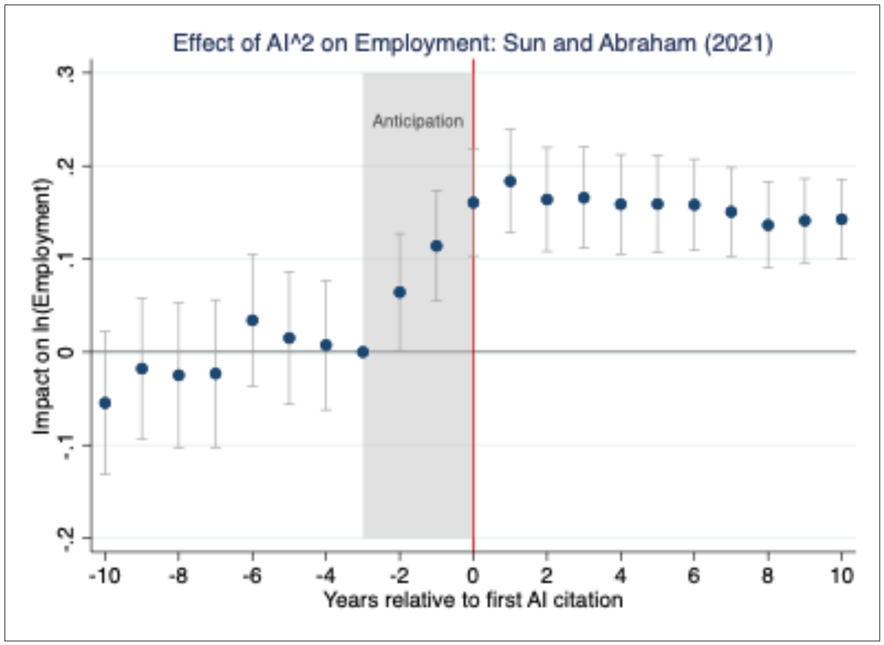

Figure 4
Two-Way Fixed Effects Analysis: Output, COGS, Profit, TFP, Capital

This figure depicts the Sun and Abraham (2021) Interaction-Weighted (IW) estimator for ten periods before and ten periods after first treatment. First treatment is the filing year of the first AI-absorbing (AI^2) patent by a firm. Using G to denote the year when a unit first becomes treated, a cohort average treatment effect on the treated, $\hat{\theta}_{g,l}$, is first estimated using the following regression

$$Y_{i,t} = \alpha_i + \lambda_t + \sum_{g} \sum_{\ell \neq -3} \theta_{g,\ell} \left(\mathbf{1}\{G_i = g\}. D_{i,t}^{\ell} \right) + \epsilon_{i,t}$$

where $D_{i,t}^{\ell} = \mathbf{1}\{t - G_i = \ell\}$ is an indicator for unit i being ℓ periods away from initial treatment at calendar time t. α_i are firm-fixed effects and λ_t are year-fixed effects. Y is one of the following: (i) the log one of Output (Figure 4.A), (ii) COGS/Sales (Figure 4.B), (iii) Profit (Figure 4.C), (iv) TFP (Figure 4.D), (v) Capital (Figure 4.E) The figures show, for each relative period ℓ, the weighted average of $\hat{\theta}_{g,l}$ using sample shares of each cohort in the relevant period ℓ as weights. The anticipatory region is shaded under the following limited anticipation assumption $E\left[Y_{i,g+\ell}^{g} - Y_{i,g+\ell}^{\infty} \middle| G_i = g\right] = 0$ for all $\ell < \delta$ where $\delta = -2$.

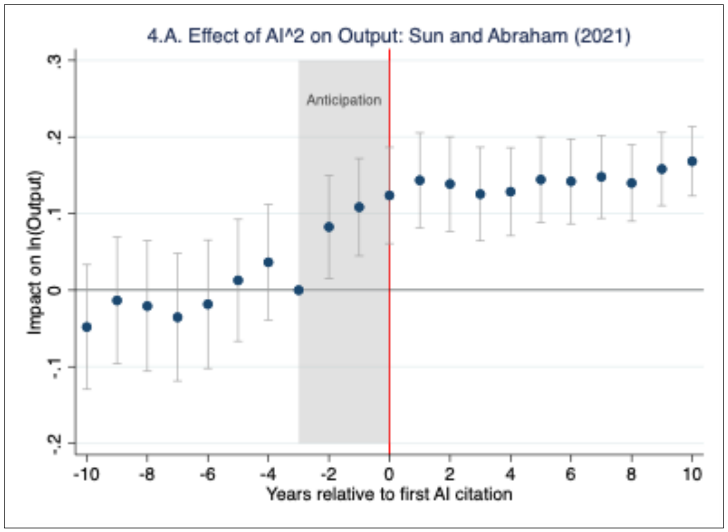

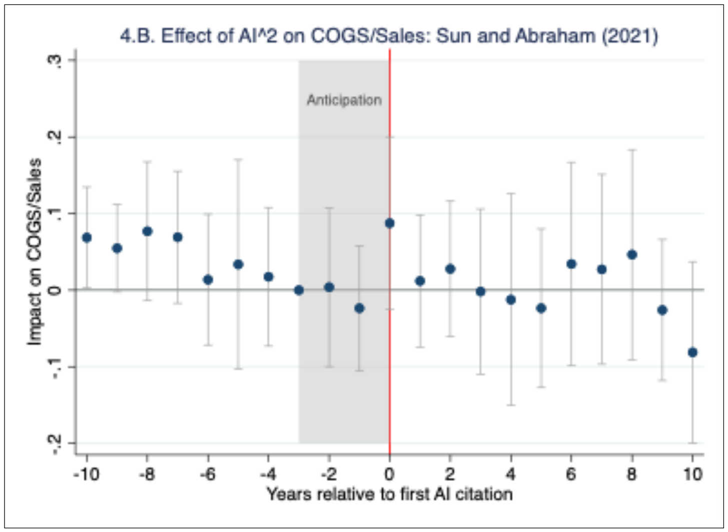

Figure 4 continued

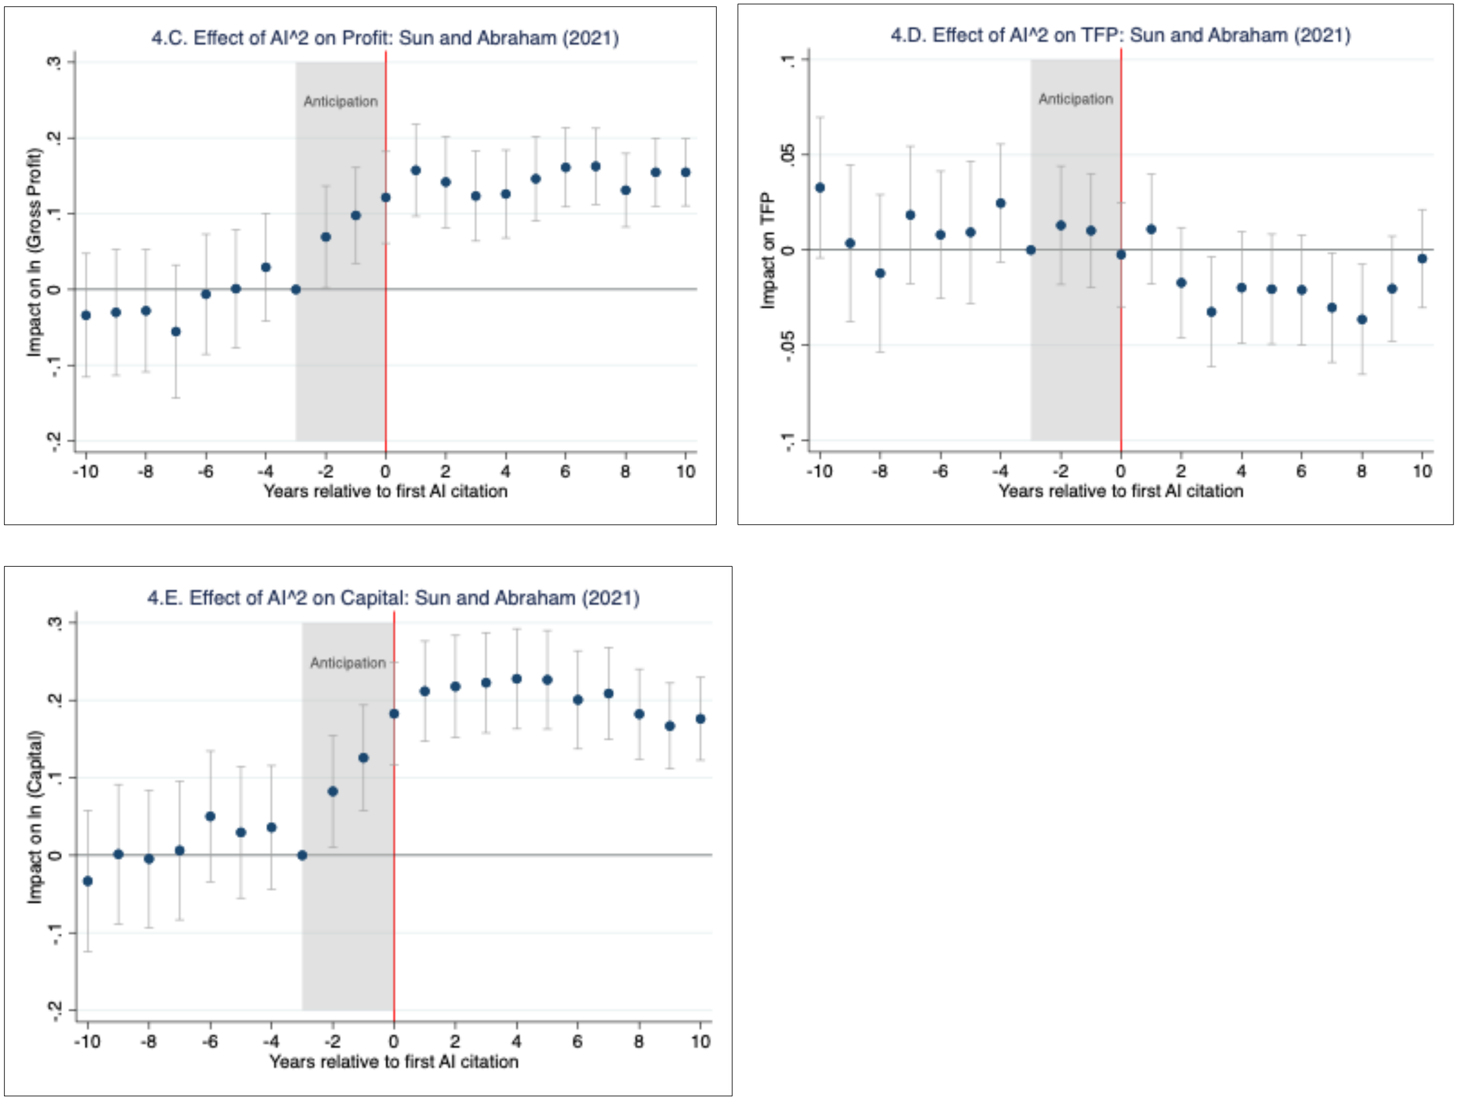

Table 1: Patent-level Summary Statistics

Panel A presents summary statistics for all utility patents issued to publicly traded firms in the U.S between 1979 and 2024. Panel B presents summary statistics for the subsample of non-IT patents only. Non-IT patents are patents *not* in the following information technology IPC classes: G06F, G06Q, G06N, G06K, G06T, H04L, H04W, H04M. $\ln(\xi)$ is the Kogan et al (2017) patent value. *Pct AI cite* is the percentage of cited patents that are AI patents. I_{AIcite} is an indicator variable taking the value 1 if the focal patent cites at least one AI patent. *Forward* captures truncation bias adjusted forward citations which involves dividing the number of forward citations a patent receives by the mean number of forward citations received by patents granted in the same year. *ln Backward* is the log of the number patents cited by the focal patent. *AI* is an indicator variable equal to 1 if focal patent is an AI patent. Patents with IPC class G06N3, G06N5, G06N7, G06N10, G06F15, and G06K9 are classified as AI patents. *ln Market cap* is calculated as of one day prior to patent grant date as the natural log of shares outstanding times share price. $ln\,(\sigma)$ is the natural log of idiosyncratic volatility in the patent grant year.

Panel A: All Patents

Variable	N	Mean	SD	p10	p25	p50	p75	p90
$\ln(\xi)$	2,533,690	0.977	2.263	-2.449	-0.168	1.445	2.509	3.463
Pct AI cite	2,533,690	0.056	0.158	0	0	0	0	0.2
I_{AIcite}	2,533,690	0.197	0.398	0	0	0	0	1
ln (1+Forward)	2,533,690	0.415	0.545	0	0	0.240	0.573	1.077
ln Backward	2,533,690	1.988	1.162	0.693	1.099	1.946	2.639	3.401
AI Focal Patent	2,533,595	0.054	0.227	0	0	0	0	0
ln Market Cap	2,345,326	16.106	2.522	12.672	14.449	16.266	18.202	19.071
ln (σ)	2,343,675	-4.087	0.467	-4.629	-4.416	-4.136	-3.811	-3.479

Panel B: Non-IT Patents

Variable	N	Mean	SD	p10	p25	p50	p75	p90
$\ln(\xi)$	1,778,737	0.813	2.290	-2.684	-0.529	1.351	2.391	3.302
Pct AI cite	1,778,737	0.012	0.067	0	0	0	0	0
I_{AIcite}	1,778,737	0.067	0.250	0	0	0	0	0
ln (1+Forward)	1,778,737	0.413	0.528	0	0.056	0.249	0.570	1.035
ln Backward	1,778,737	2.031	1.133	0.693	1.386	1.946	2.639	3.401
AI Focal Patent	1,778,737	0	0	0	0	0	0	0
ln Market Cap	1,633,865	15.631	2.446	12.290	14.024	15.796	17.541	18.647
ln (σ)	1,632,719	-4.063	0.473	-4.622	-4.394	-4.113	-3.792	-3.450

Table 2: AI-Absorbing Innovation and Patent Value

This table presents results from estimating the following patent-level equation:

$$ln(\xi_p) = \alpha + \beta\ PctAIcite_p + \gamma_1 X_i + \gamma_2 \ln Backward_p + \mu_i + \rho_{k,t} + \epsilon_p$$

where ξ is the Kogan et al (2017) value of patent p, *PctAIcite* is the share of a patent's backward citations that cite an AI patent, *ln Backward* is the log of the number patents cited by the patent, X is a vector of firm-level control variables including log market capitalization and log idiosyncratic volatility, $\mu_i$are firm-fixed effects, $\rho_{k,t}$ are technology class times grant year fixed effects. Standard errors are clustered by year. ***, **,* indicate significance at the 1%, 5%, and 10% respectively. All variable definitions are provided in Appendix A. The sample period is 1979 to 2024.

Dependent Variable:		$ln(\xi)$	
Sample:	All Patents	Non-IT Patents	IT Patents
	(1)	(2)	(3)
PctAIcite	-0.005	0.022***	-0.003
	(-1.037)	(2.803)	(-0.657)
ln Market Cap	0.893***	0.906***	0.844***
	(75.487)	(71.692)	(57.806)
ln (σ)	0.755***	0.761***	0.723***
	(29.152)	(26.265)	(19.610)
ln Backward	0.005***	0.007***	0.001
	(4.400)	(4.652)	(0.519)
Observations	2,336,260	1,625,472	708,417
R-squared	0.937	0.939	0.932
Firm FE	Yes	Yes	Yes
PatentClass*Year FE	Yes	Yes	Yes
Clustered SE	Year	Year	Year

Table 3: Firm-level Summary Statistics

Summary statistics for firms excluding the software and communications sector over the sample period 1979 to 2024. Panel A includes all firms. Panel B is restricted to innovative firm-years defined as firm-years with at least one patent in the period *t* through *t-5*. All variable definitions are provided in Appendix A.

Panel A: All firms

Variable	N	Mean	SD	p10	p25	p50	p75	p90
AI^2	264,962	0.026	0.132	0	0	0	0	0
I_{AI^2}	264,962	0.066	0.249	0	0	0	0	0
ln Output	203,349	4.654	2.442	1.650	3.096	4.723	6.312	7.744
ln Profit	212,161	3.672	2.240	0.890	2.183	3.668	5.180	6.542
COGS/Sales	204,334	0.982	3.605	0.280	0.473	0.652	0.793	0.917
TFP	100,185	-0.328	0.475	-0.807	-0.531	-0.317	-0.102	0.163
ln Employment	215,887	-0.167	2.340	-3.194	-1.814	-0.124	1.504	2.879
ln Capital Stock	220,498	3.750	2.805	0.235	1.813	3.649	5.707	7.510
$ln\ (\sigma)$	250,656	-3.615	0.625	-4.409	-4.082	-3.649	-3.183	-2.778

Panel B: Innovative subsample

Variable	N	Mean	SD	p10	p25	p50	p75	p90
AI^2	89,999	0.075	0.217	0	0	0	0	0.25
I_{AI^2}	89,999	0.196	0.397	0	0	0	0	1
ln Output	73,441	5.023	2.685	1.618	3.283	5.123	6.969	8.451
ln Profit	70,681	4.275	2.355	1.293	2.664	4.281	5.969	7.404
COGS/Sales	73,721	1.297	4.756	0.327	0.493	0.652	0.785	0.947
TFP	47,293	-0.304	0.496	-0.794	-0.506	-0.286	-0.068	0.207
ln Employment	78,614	0.258	2.369	-2.847	-1.556	0.273	2.054	3.444
ln Capital Stock	79,444	4.056	2.950	0.337	1.951	3.929	6.204	8.059
$ln\ (\sigma)$	80,413	-3.555	0.592	-4.315	-4.002	-3.576	-3.136	-2.771

Table 4: AI^2 and Employment

This table presents the relation between AI-absorbing innovation (AI^2) and firm-level outcomes. We present the coefficients β_τ β_τ for cumulative forward time periods $\tau \in \{1, \ldots, 5\}$ from estimations of the following model:

$$\ln(Emp_{i,t+\tau}) - \ln(Emp_{i,t}) = \alpha_0 + \beta_\tau AI^\wedge 2_{it} + \gamma_1 X'_{it} + \delta_{jt} + \epsilon_{it}$$

where Emp is number of employees, AI^2 is fraction of the firm's patents that are AI-citing, X is a vector of firm-level control variables including log value of capital stock, log idiosyncratic volatility, ln(1+$Patents_{i,t-1}$), and log ($Emp_{i,t-1}$), $\delta_{j,t}$ are SIC3 times year fixed effects. Standard errors are double-clustered by firm and year. ***, **,* indicate significance at the 1%, 5%, and 10% respectively. All variable definitions are provided in Appendix A. The sample period is 1979 to 2024. Panel A is the full sample of publicly traded firms, whereas Panel B restricts firms serving as comparable to innovative firm-years defined as firm-years with at least one patent in the period *t* through *t-5*.

	Dependent Variable: Change in Log Employment over Horizon:				
	t+1	t+2	t+3	t+4	t+5
Panel A: Full Panel of Innovative and Non-Innovative Comparables					
AI^2	0.029***	0.046***	0.075***	0.080***	0.089***
	(4.648)	(3.715)	(4.775)	(3.516)	(3.266)
ln (1+Patents) (lag)	0.006***	0.012***	0.017***	0.023***	0.026***
	(5.330)	(6.158)	(6.250)	(6.478)	(6.330)
ln Employment (lag)	-0.043***	-0.074***	-0.097***	-0.118***	-0.138***
	(-17.973)	(-19.536)	(-19.244)	(-18.340)	(-17.330)
Observations	163,475	147,278	133,031	120,545	109,560
R-squared	0.105	0.128	0.141	0.154	0.167
Panel B: Panel Restricted to Innovative Comparables					
AI^2	0.026***	0.043***	0.072***	0.076***	0.083***
	(3.855)	(3.171)	(4.371)	(3.186)	(2.954)
ln (1+Patents) (lag)	0.009***	0.016***	0.022***	0.028***	0.031***
	(6.097)	(6.488)	(6.900)	(6.927)	(6.367)
ln Employment (lag)	-0.049***	-0.079***	-0.099***	-0.116***	-0.136***
	(-15.768)	(-14.717)	(-13.726)	(-11.935)	(-11.475)
Observations	64,229	58,976	54,079	49,564	45,467
R-squared	0.132	0.151	0.159	0.167	0.175
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 5: Instrumented AI^2 and Employment

This table presents an instrumental variable analysis of the relation between AI-absorbing innovation (AI^2) and firm-level outcomes. We present the coefficients β_τ for cumulative forward time periods $\tau \in \{1, \ldots, 5\}$ from estimations of the following model:

$$\ln(Emp_{i,t+\tau}) - \ln(Emp_{i,t}) = \alpha_0 + \beta_\tau \widehat{AI^2}_{i,t} + \gamma_1 X'_{i,t} + +\delta_{jt} + \psi_{c,t} + \epsilon_{i,t}$$

where Emp is number of employees. $\widehat{\text{AI}^\wedge 2}$ is the fitted value from a first stage regression of AI^2 on *Laboratory AI*, the number of AI patents granted to universities and federal labs located in the firm's commuting zone. X is a vector of firm-level control variables including log value of capital stock, log idiosyncratic volatility, ln(1+$\text{Patents}_{i,t-1}$), and log ($\text{Emp}_{i,t-1}$). $\delta_{j,t}$ are SIC3 times year fixed effects and $\psi_{c,t}$ are commuting zone times year fixed effects. Standard errors are double-clustered by firm and year. ***, **,* indicate significance at the 1%, 5%, and 10% respectively. All variable definitions are provided in Appendix A. The sample period is 1979 to 2024. Panel A is the full sample of publicly traded firms, whereas Panel B restricts firms serving as comparable to innovative firm-years defined as firm-years with at least one patent in the period *t* through *t-5*.

	t+1	t+2	t+3	t+4	t+5
Panel A: Full Panel of Innovative and Non-Innovative Comparables					
First Stage					
Laboratory AI First stage F-statistic	321.57	299.04	279.10	251.84	247.99
Second Stage					
$\widehat{AI^\wedge 2}$	0.151***	0.148***	0.167**	0.209**	0.236**
	(4.88)	(2.96)	(2.20)	(2.33)	(2.14)
ln Employment (lag)	-0.046***	-0.078***	-0.103***	-0.126***	-0.147***
	(-17.47)	(-19.81)	(-19.06)	(-18.30)	(-17.26)
Observations	139473	125846	113727	103014	93562
Panel B: Panel Restricted to Innovative Comparables					
First Stage					
Laboratory AI First stage F-statistic	274.25	242.68	225.33	197.73	193.18
Second Stage					
$\widehat{AI^\wedge 2}$	0.218***	0.193**	0.192*	0.274**	0.248
	(5.59)	(2.60)	(1.72)	(2.18)	(1.53)
ln Employment (lag)	-0.051***	-0.085***	-0.104***	-0.123***	-0.145***
	(-14.70)	(-14.13)	(-13.92)	(-12.21)	(-11.71)
Observations	53179	48730	44561	40754	37262
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Commuting Zone x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 6: AI^2 and Firm Performance Mechanisms

This table presents an instrumental variable analysis of the relation between AI-absorbing innovation (AI^2) and firm-level outcomes. We present the coefficients β_τ for cumulative forward time periods $\tau \in \{1, \dots, 5\}$ from estimations of the following model:

$$\ln(Y_{i,t+\tau}) - \ln(Y_{it}) = \alpha_0 + \beta_\tau AI^\wedge 2_{it} + \gamma_1 X'_{it} + \delta_{jt} + \epsilon_{it}$$

where Y are firm outcomes represented in the Panel names; AI^2 is the fraction of the firm's patents that are AI-citing; X is a vector of firm-time level control variables including log value of capital stock, the log number of employees, log idiosyncratic volatility, ln(1+$Patents_{i,t-1}$), and ln($Y_{i,t-1}$); δ_{jt} are SIC3 times year fixed effects. Standard errors are double-clustered by firm and year. ***, **, * indicate significance at the 1%, 5%, and 10% respectively. All variable definitions are provided in Appendix A. The sample period is 1979 to 2024. Panel A is the full sample of publicly traded firms. Panel B restricts firms serving as comparable to innovative firm-years defined as firm-years with at least one patent in the period *t* through *t-5*

	t+1	t+2	t+3	t+4	t+5
Panel A: Full Panel of Innovative and Non-Innovative Comparables					
Panel A.1: Output					
AI^2	0.032***	0.058***	0.080***	0.103***	0.114***
	(5.304)	(3.723)	(3.852)	(4.138)	(3.566)
Observations	145,496	131,068	118,961	108,209	98,595
R-squared	0.116	0.149	0.157	0.166	0.179
Panel A.2: COGS/Sales					
AI^2	-0.010	0.039	0.012	-0.125***	-0.051
	(-0.422)	(0.514)	(0.155)	(-2.776)	(-0.868)
Observations	146,891	132,277	120,004	109,101	99,380
R-squared	0.028	0.041	0.040	0.057	0.051
Panel A.3: Profit					
AI^2	0.048***	0.082***	0.104***	0.122***	0.136***
	(5.743)	(5.554)	(5.908)	(5.420)	(4.885)
Observations	150,067	135,850	123,343	112,283	102,445
R-squared	0.148	0.174	0.184	0.193	0.208
Panel A.4: TFP					
AI^2	0.021**	0.027*	0.038*	0.039*	0.031
	(2.208)	(1.782)	(2.007)	(1.703)	(1.159)
Observations	76,732	68,818	62,327	56,641	51,606
R-squared	0.097	0.115	0.138	0.162	0.176
Panel A.5: Capital					
AI^2	0.032***	0.045**	0.072***	0.075**	0.072**
	(3.694)	(2.581)	(3.166)	(2.461)	(2.031)
Observations	166,305	150,024	135,611	122,943	111,731
R-squared	0.122	0.148	0.166	0.178	0.188
Industry x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 6 (continued)

	t+1	t+2	t+3	t+4	t+5
	Panel B: Panel Restricted to Innovative Comparables				
	Panel B.1: Output				
AI^2	0.029***	0.056***	0.078***	0.097***	0.107***
	(4.226)	(3.101)	(3.296)	(3.517)	(3.198)
Observations	58,671	53,823	49,538	45,570	41,917
R-squared	0.112	0.149	0.159	0.171	0.181
	Panel B.2: COGS/Sales				
AI^2	-0.020	0.029	-0.014	-0.191***	-0.124
	(-0.698)	(0.317)	(-0.148)	(-2.901)	(-1.493)
Observations	59,120	54,212	49,876	45,850	42,163
R-squared	0.028	0.041	0.038	0.060	0.051
	PanelBA.3: Profit				
AI^2	0.039***	0.078***	0.103***	0.117***	0.135***
	(4.377)	(4.429)	(5.355)	(4.793)	(4.612)
Observations	56,495	52,147	48,111	44,336	40,847
R-squared	0.181	0.208	0.223	0.226	0.235
	Panel B.4: TFP				
AI^2	0.022**	0.031*	0.038*	0.038	0.028
	(2.237)	(1.944)	(1.905)	(1.583)	(0.977)
Observations	38,138	34,866	32,061	29,484	27,135
R-squared	0.121	0.141	0.164	0.187	0.196
	Panel B.5: Capital				
AI^2	0.027***	0.039*	0.066**	0.065*	0.054
	(3.000)	(1.987)	(2.609)	(1.939)	(1.411)
Observations	64,904	59,634	54,698	50,129	45,987
R-squared	0.143	0.164	0.181	0.192	0.199
Industry x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 7: Instrumented AI^2 and Firm Performance Mechanisms

This table presents the relation between AI-absorbing innovation (AI^2) and firm-level outcomes. We present the coefficients β_5 for cumulative forward time periods $\tau = 5$ from estimations of the following model:

$$\ln(Y_{i,t+5}) - \ln(Y_{it}) = \alpha_0 + \beta_5 \widehat{AI\text{^}2}_{it} + \gamma_1 X_{it}' + \delta_{jt} + \epsilon_{it}$$

where Y are firm outcomes represented in the Panel names; $\widehat{\text{AI^2}}$ is the fitted value from a first stage regression of AI^2 on the number of AI patents granted to universities and federal labs located in the firm's commuting zone; X is a vector of firm-time level control variables including log value of capital stock, the log number of employees, log idiosyncratic volatility, ln(1+$\text{Patents}_{i,t-1}$), and ln($Y_{i,t-1}$); δ_{jt} are SIC3 times year fixed effects. Standard errors are double-clustered by firm and year. ***, **,* indicate significance at the 1%, 5%, and 10% respectively. All variable definitions are provided in Appendix A. The sample period is 1979 to 2024.

	t+1	t+2	t+3	t+4	t+5
Panel A: Full Panel of Innovative and Non-Innovative Comparables					
Panel A.1: Output					
$\widehat{AI\text{^}2}$	0.121	0.259**	0.171	0.255*	0.330**
	(1.61)	(2.68)	(1.51)	(1.96)	(2.37)
Observations	123909	111701	101402	92198	83890
First stage F-statistic	327.33	305.12	277.96	251.33	243.30
Panel A.2: COGS/Sales					
$\widehat{AI\text{^}2}$	-0.338	-0.084	-0.053	-0.209	-0.114
	(-0.54)	(-0.18)	(-0.11)	(-0.48)	(-0.19)
Observations	125082	112702	102263	92935	84535
First stage F-statistic	341.43	313.77	283.54	251.74	243.05
Panel A.3: Profit					
$\widehat{AI\text{^}2}$	0.182***	0.312***	0.300***	0.344***	0.374***
	(5.02)	(5.37)	(4.00)	(3.95)	(4.05)
Observations	127541	115561	104924	95473	87000
First stage F-statistic	368.89	343.88	315.03	286.49	276.71
Panel A.4: TFP					
$\widehat{AI\text{^}2}$	0.001	0.059	0.087	0.083	0.097
	(0.03)	(0.98)	(1.00)	(1.01)	(1.11)
Observations	64881	58196	52680	47806	43442
First stage F-statistic	214.81	215.49	201.00	178.23	162.91
Panel A.5: Capital					
$\widehat{AI\text{^}2}$	0.151***	0.148***	0.167**	0.209**	0.236**
	(4.88)	(2.96)	(2.20)	(2.33)	(2.14)
Observations	139473	125846	113727	103014	93562
First stage F-statistic	321.57	299.04	279.10	251.84	247.99
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Commuting Zone x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 7 (continued)

	t+1	t+2	t+3	t+4	t+5
	Panel B: Panel Restricted to Innovative Comparables				
	Panel B.1: Output				
$\widehat{AI}$^2	0.170*	0.306**	0.113	0.129	0.142
	(1.68)	(2.27)	(0.68)	(0.67)	(0.74)
Observations	48212	44123	40489	37145	34060
First stage F-statistic	265.11	232.34	209.13	184.93	181.55
	Panel B.2: COGS/Sales				
$\widehat{AI}$^2	-0.322	0.290	0.061	-0.321	0.173
	(-0.36)	(0.36)	(0.07)	(-0.52)	(0.23)
Observations	48609	44466	40796	37399	34276
First stage F-statistic	279.19	239.02	214.78	185.71	181.31
	PanelBA.3: Profit				
$\widehat{AI}$^2	0.228***	0.363***	0.327***	0.344***	0.360***
	(4.63)	(4.65)	(3.43)	(3.30)	(2.94)
Observations	46029	42391	38983	35829	32908
First stage F-statistic	304.50	264.93	240.61	217.70	210.50
	Panel B.4: TFP				
$\widehat{AI}$^2	0.006	0.122	0.105	0.120	0.130
	(0.11)	(1.47)	(0.93)	(1.04)	(1.02)
Observations	30585	27880	25561	23428	21452
First stage F-statistic	168.39	162.85	149.59	139.37	125.61
	Panel B.5: Capital				
$\widehat{AI}$^2	0.422***	0.662***	0.609***	0.562**	0.412*
	(5.94)	(4.77)	(3.37)	(2.61)	(1.84)
Observations	53617	49166	44969	41106	37581
First stage F-statistic	276.30	243.69	217.15	196.24	195.10
Industry x Year FE	Yes	Yes	Yes	Yes	Yes
Commuting Zone x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 8: Automation and Process Patents Statistics

Presented are statistics of patent-level (restricted to non-IT patent sample) and firm-level variables (restricted to non-IT firms only) for 1979-2024. I_{AIcite} is an indicator variable taking the value 1 if the focal patent cites at least one AI patent; $I_{Automation}$ is an indicator variable taking the value 1 if the focal patent is an automation patent; $I_{Process}$ is an indicator variable that takes the value 1 if the more than 50% of the focal patent's claims are process claims. Firm-level variables are: *AI^2* is fraction of the firm's patents that are AI-citing. *Automation* is the fraction of a firm's patents that are automation patents. *Process* is the fraction of a firm's patents with more than 50% process claims. See Appendix C for the classification of process claims, automation patents, and non-IT patents. Panels B and C present correlational estimates to assess partial R-squared among patent taxonomy at the patent and firm level respectively. ***, **,* indicate significance at the 1%, 5%, and 10% respectively. Variable definitions are provided in Appendix A.

Panel A: Summary statistics

Variable	N	Mean	SD	p10	p25	p50	p75	p90
Patent level								
I_{AIcite}	1,778,737	0.067	0.250	0	0	0	0	0
$I_{Automation}$	1,777,127	0.139	0.346	0	0	0	0	1
$I_{Process}$	1,778,724	0.285	0.452	0	0	0	1	1
Firm level								
AI^2	264,962	0.026	0.132	0	0	0	0	0
Automation	264,962	0.045	0.169	0	0	0	0	0.053
Process	264,962	0.072	0.213	0	0	0	0	0.281

Panel B: Patent-Level Distinction among AI Absorbing, Automation, and Process Patents

	(1)	(2)	(3)	(4)
	Dependent Variable: Indicator for AI Cite (I_{AIcite})			
$I_{Automation}$	0.118***	0.058***		
	(36.37)	(38.99)		
$I_{Process}$			0.024***	0.020***
			(11.68)	(26.25)
Observations	1,777,221	1,769,666	1,778,819	1,771,262
Partial R-sq	0.027	0.007	0.002	0.001
Fixed effects	No	Firm, Class x Year	No	Firm, Class x Year
Clustering	Year	Year	Year	Year

Panel C: Firm-Level Distinction among AI Absorbing, Automation, and Process Patents

	(1)	(2)	(3)	(4)
	Dependent Variable: AI^2			
Automation	0.304***	0.278***		
	(0.013)	(0.013)		
Process			0.133***	0.126***
			(0.011)	(0.010)
Observations	238,197	232,730	238,197	232,730
Partial R-sq	0.155	0.140	0.043	0.041
Fixed effects	No	SIC3 x Year	No	SIC3 x Year
Control for # patents	Yes	Yes	Yes	Yes
Clustering	Firm, year	Firm, year	Firm, year	Firm, year

Table 9: Product and Process Innovation as AI^2

This table presents an instrumental variable analysis of the relation between AI-absorbing innovation (AI^2) and firm-level outcomes. We present the coefficients β_τ for cumulative forward time periods $\tau \in \{1, \ldots, 5\}$ from estimations of the following model:

$$\ln(Y_{i,t+\tau}) - \ln(Y_{it}) = \alpha_0 + \beta_{1\tau} Product\ AI\text{^}2_{it} + \beta_{2\tau} Process\ AI\text{^}2_{it} + \beta_{3\tau} Process_{it} + \gamma_1 X'_{it} + \delta_{jt} + \epsilon_{it}$$

where Y are firm outcomes represented in the Panel names; $Product\ AI$^2 is the fraction of the firm's patents that are AI-citing product patents; $Process\ AI$^2 is the fraction of the firm's patents that are AI-citing process patents. $Process$ is the fraction of a firm's patents that are process patents. Process patents are patents with more than 50% process claims. The rest are classified as product patents. The classification of claims is described in Appendix C. X is a vector of firm-time level control variables including log value of capital stock, the log number of employees, log idiosyncratic volatility, $\ln(1+Patents_{i,t-1})$, and $\ln(Y_{i,t-1})$; δ_{jt} are SIC3 times year fixed effects. Standard errors are double-clustered by firm and year. ***, **,* indicate significance at the 1%, 5%, and 10% respectively. All variable definitions are provided in Appendix A. The sample period is 1979 to 2024.

	t+1	t+2	t+3	t+4	t+5
			Panel A: Employment		
Product AI^2	0.027***	0.040**	0.048**	0.060**	0.054*
	(2.85)	(2.54)	(2.50)	(2.40)	(1.76)
Process AI^2	0.025*	0.052**	0.116***	0.111***	0.149***
	(1.87)	(2.04)	(3.48)	(2.86)	(3.21)
Process	0.010**	0.003	-0.002	-0.003	-0.015
	(2.16)	(0.48)	(-0.21)	(-0.26)	(-0.98)
Observations	163475	147278	133031	120545	109560
R-squared	0.105	0.128	0.141	0.154	0.167
			Panel B: Output		
Product AI^2	0.052***	0.052***	0.069***	0.089***	0.097***
	(5.14)	(2.84)	(2.98)	(3.45)	(2.87)
Process AI^2	-0.004	0.047	0.078**	0.115***	0.121**
	(-0.16)	(1.45)	(2.09)	(2.79)	(2.32)
Process	0.011	0.026**	0.021	0.009	0.021
	(1.07)	(2.13)	(1.27)	(0.44)	(0.87)
Observations	145496	131068	118961	108209	98595
R-squared	0.117	0.149	0.157	0.166	0.179
			Panel C: COGS/Sale		
Product AI^2	-0.092*	0.060	-0.041	-0.141**	-0.047
	(-1.79)	(0.53)	(-0.40)	(-2.05)	(-0.70)
Process AI^2	0.067	-0.056	0.121	-0.170	-0.130
	(0.43)	(-0.48)	(0.80)	(-1.10)	(-0.81)
Process	0.048	0.083	-0.044	0.084	0.090
	(0.48)	(0.84)	(-0.50)	(0.78)	(0.87)
Observations	146891	132277	120004	109101	99380
R-squared	0.028	0.041	0.040	0.057	0.051

Table 9 (continued)

Panel D: Profit					
Product AI^2	0.056***	0.089***	0.092***	0.124***	0.133***
	(5.73)	(5.23)	(3.92)	(4.77)	(4.24)
Process AI^2	0.021	0.048	0.101***	0.100**	0.125**
	(1.21)	(1.67)	(2.81)	(2.49)	(2.70)
Process	0.017**	0.028**	0.025*	0.024	0.020
	(2.09)	(2.45)	(1.81)	(1.31)	(1.06)
Observations	150067	135850	123343	112283	102445
R-squared	0.148	0.174	0.184	0.193	0.208
Panel E: TFP					
Product AI^2	0.009	-0.009	0.017	0.024	0.030
	(0.92)	(-0.53)	(0.82)	(0.88)	(1.04)
Process AI^2	0.043**	0.086***	0.081**	0.082**	0.064
	(2.08)	(3.39)	(2.33)	(2.42)	(1.31)
Process	-0.005	-0.005	-0.015	-0.026**	-0.038***
	(-0.86)	(-0.55)	(-1.57)	(-2.27)	(-2.78)
Observations	76732	68818	62327	56641	51606
R-squared	0.097	0.115	0.138	0.162	0.176
Panel F: Capital					
Product AI^2	0.015	0.036*	0.052**	0.051	0.036
	(1.22)	(1.81)	(2.15)	(1.60)	(0.92)
Process AI^2	0.040**	0.030	0.077*	0.073	0.095
	(2.16)	(0.94)	(1.81)	(1.45)	(1.66)
Process	0.018**	0.035***	0.029*	0.043***	0.032
	(2.44)	(3.54)	(1.92)	(2.72)	(1.58)
Observations	166305	150024	135611	122943	111731
R-squared	0.122	0.148	0.166	0.178	0.188
Industry x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 10: The Role of AI^2 Automation and Routine Automation

This table presents an instrumental variable analysis of the relation between AI-absorbing innovation (AI^2) and firm-level outcomes. We present the coefficients β_τ for cumulative forward time periods $\tau \in \{1, \ldots, 5\}$ from estimations of the following model:

$$\ln(Y_{i,t+\tau}) - \ln(Y_{it}) = \alpha_0 + \beta_{1\tau} Automation\ AI\text{^}2_{it} + \beta_{2\tau} NonAutomation\ AI\text{^}2_{it} + \beta_{3\tau} Automation_{it} + \gamma_1 X'_{it} + \delta_{jt} + \epsilon_{it}$$

where Y are firm outcomes represented in the Panel names; $Automation\ AI\text{^}2$ is the fraction of the firm's patents that are AI-citing automation patents; $NonAutomation\ AI\text{^}2$ is the fraction of the firm's patents that are AI-citing non-automation patents. $Automation$ is the fraction of a firm's patents that are automation patents. The classification of automation patents is described in Appendix C. X is a vector of firm-time level control variables including log value of capital stock, the log number of employees, log idiosyncratic volatility, ln(1+$Patents_{i,t-1}$), and ln($Y_{i,t-1}$); δ_{jt} are SIC3 times year fixed effects. Standard errors are double-clustered by firm and year. ***, **,* indicate significance at the 1%, 5%, and 10% respectively. All variable definitions are provided in Appendix A. The sample period is 1979 to 2024. Panel A is the full sample of publicly traded firms. Panel B restricts firms serving as comparable to innovative firm-years defined as firm-years with at least one patent in the period *t* through *t-5*

	t+1	t+2	t+3	t+4	t+5
			Panel A: Employment		
Automation AI^2	0.009	0.027	0.067**	0.056	0.086**
	(0.84)	(1.66)	(2.52)	(1.64)	(2.28)
Non-Automation AI^2	0.030***	0.034*	0.052**	0.063*	0.056
	(3.25)	(1.98)	(2.35)	(1.98)	(1.49)
Automation	0.021***	0.031***	0.031**	0.040**	0.033
	(3.43)	(2.88)	(2.18)	(2.14)	(1.50)
Observations	163475	147278	133031	120545	109560
R-squared	0.105	0.128	0.141	0.154	0.167
			Panel B: Output		
Automation AI^2	0.006	0.021	0.055*	0.063	0.092*
	(0.40)	(0.81)	(1.99)	(1.64)	(1.71)
Non-Automation AI^2	0.036***	0.064***	0.080**	0.116***	0.099**
	(2.82)	(2.85)	(2.70)	(3.23)	(2.36)
Automation	0.023*	0.033**	0.025	0.028	0.037
	(1.70)	(2.27)	(1.15)	(1.26)	(1.21)
Observations	145496	131068	118961	108209	98595
R-squared	0.117	0.149	0.157	0.166	0.179
			Panel C: COGS/Sale		
Automation AI^2	0.007	0.078	-0.140	-0.210	-0.210
	(0.06)	(0.54)	(-1.18)	(-1.30)	(-1.37)
Non-Automation AI^2	0.021	-0.028	0.050	-0.179*	-0.016
	(0.23)	(-0.32)	(0.35)	(-1.84)	(-0.17)
Automation	-0.046	0.023	0.117	0.136	0.125
	(-0.36)	(0.38)	(0.97)	(0.93)	(0.92)
Observations	146891	132277	120004	109101	99380
R-squared	0.028	0.041	0.040	0.057	0.051

Table 10 (continued)

			Panel D: Profit		
Automation AI^2	0.029**	0.030	0.048*	0.057*	0.079*
	(2.56)	(1.61)	(1.85)	(1.71)	(1.92)
Non-Automation AI^2	0.045***	0.094***	0.119***	0.128***	0.139***
	(3.39)	(4.52)	(4.38)	(4.01)	(3.40)
Automation	0.021**	0.039***	0.042**	0.059***	0.055**
	(2.58)	(3.00)	(2.20)	(2.94)	(2.33)
Observations	150067	135850	123343	112283	102445
R-squared	0.148	0.174	0.184	0.193	0.208
			Panel E: TFP		
Automation AI^2	0.017	0.007	0.041	0.023	0.015
	(1.06)	(0.34)	(1.37)	(0.64)	(0.36)
Non-Automation AI^2	0.021	0.035	0.041	0.040	0.031
	(1.49)	(1.58)	(1.35)	(1.36)	(0.91)
Automation	0.003	0.012	-0.006	0.014	0.017
	(0.42)	(1.20)	(-0.43)	(0.89)	(0.90)
Observations	76732	68818	62327	56641	51606
R-squared	0.097	0.115	0.138	0.162	0.176
			Panel F: Capital		
Automation AI^2	0.013	-0.014	0.043	0.037	0.083
	(0.73)	(-0.52)	(1.16)	(0.77)	(1.53)
Non-Automation AI^2	0.027**	0.047**	0.056*	0.069*	0.036
	(2.36)	(2.13)	(1.93)	(1.74)	(0.76)
Automation	0.023**	0.058***	0.044*	0.044	0.024
	(2.20)	(3.50)	(1.92)	(1.61)	(0.74)
Observations	166305	150024	135611	122943	111731
R-squared	0.122	0.148	0.166	0.178	0.188
Industry x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

APPENDIX A

Table A1. Variable Description

Variable	Description
ξ	Kogan et al (2017) value of a patent
σ	Idiosyncratic volatility in the patent grant year
AI^2	The percentage of patents in a firm-year that cite an AI patent
AI Focal Patent	Indicator variable equal to 1 if focal patent is an AI patent identified as a patent with IPC classification G06N3, G06N5, G06N7, G06N10, G06F15, or G06K9. See Table A2 for a description of these classifications.
Automation	Fraction of a firm's patents that are automation patents. Appendix C.1. describes the classification of automation patents
Backward	Natural log of the number patents cited by the focal patent (i.e. backward citations)
Capital stock	COMPUTAT: ppent, deflated by the NIPA price of equipment
COGS/Sales	COMPUTAT: cogs/(sale+change in inventory)
Forward	Truncation bias adjusted forward citations calculated by dividing the number of forward citations a patent receives by the mean number of forward citations received by patents granted in the same year
Employment (Emp)	Number of employees
Exogenous AI	Number of AI patents granted to universities and federal labs located in the firm's commuting zone
I_{AI^2}	Indicator variable equal to one for firm-years with at least one patent that cites an AI patent, and zero otherwise
I_{AIcite}	Indicator variable taking the value 1 if the patent cites at least one AI patent. AI patents are patents with IPC classification G06N3, G06N5, G06N7, G06N10, G06F15, or G06K9. See Table A2 for a description of these classifications.
$I_{Automation}$	Indicator variable equal to 1 if focal patent is an automation patent identified based on textual analysis of patent description. Appendix C.1. describes the classification of automation patents.
$I_{Process}$	Indicator variable that takes the value 1 if the more than 50% of the focal patent's claims are process claims. Appendix C.2. describes the classification of process claims.
Patents	Number of patents
Market cap	Shares outstanding times share price one day prior to patent grant date
Output	COMPUSTAT: sale plus change in inventories (invt) deflated by the CPI
PctAIcite	Percentage of cited patents that are AI patents
Process	Fraction of a firm's patents with more than 50% process claims. Appendix C.2. describes the classification of process claims.
Profit	Gross profit calculated as sales minus cost of goods sold
TFP	Total factor productivity from İmrohoroğlu and Tüzel (2014)

Table A2. Description of Artificial Intelligence (AI) IPC Classes

IPC Class	Description
G06N3	Computing arrangements based on biological models
G06N5	Computing arrangements using knowledge-based models
G06N7	Computing arrangements based on specific mathematical models
G06N10	Quantum computing, i.e. information processing based on quantum-mechanical phenomena.
G06F15	Digital computers in general; Data processing equipment
G06K9	Methods or arrangements for recognizing patterns

Table A3. Description of Information Technology (IT) IPC Classes

IPC Class	Description
G06F	Electric digital data processing
G06Q	Data processing systems or methods, specially adapted for administrative, commercial, financial, managerial, supervisory or forecasting purposes.
G06N	Computing arrangements based on specific computational models
G06K	Recognition of data; Presentation of data; Record carriers; Handling record carriers
G06T	Image data processing or generation, in general
H04L	Transmission of digital information, e.g. telegraphic communication
H04W	Wireless communication networks
H04M	Telephonic communications

APPENDIX B: Patent-level Robustness Tests

Table B1: AI-Absorbing Innovation and Patent Value (using I_{AIcite})

This table presents results from estimating the following patent-level equation:

$$ln(\xi_p) = \alpha + \beta\, I_{AIcite_p} + \gamma_1 X_i + \gamma_2 \ln Backward_p + \mu_i + \rho_{k,t} + \epsilon_p$$

where ξ is the Kogan et al (2017) value of patent p, I_{AIcite} is the share of a patent's backward citations that cite an AI patent, *ln Backward* is the log of the number patents cited by the patent, X is a vector of firm-level control variables including log market capitalization and log idiosyncratic volatility, μ_i are firm-fixed effects, $\rho_{k,t}$ are technology class times grant year fixed effects. Standard errors are clustered by year. ***, **,* indicate significance at the 1%, 5%, and 10% respectively. All variable definitions are provided in Appendix A. The sample period is 1979 to 2024.

Dependent Variable in all columns: $ln(\xi)$

	All patents		Non-IT patents		IT patents	
I_{AIcite}	0.004**	-0.001	0.011***	0.005*	0.000	-0.000
	(2.408)	(-0.653)	(3.105)	(1.689)	(0.203)	(-0.244)
Market cap (logs)	0.893***	0.893***	0.906***	0.906***	0.844***	0.844***
	(75.404)	(75.478)	(71.604)	(71.691)	(57.830)	(57.815)
ln (σ)	0.755***	0.755***	0.761***	0.761***	0.723***	0.723***
	(29.144)	(29.153)	(26.238)	(26.267)	(19.610)	(19.610)
Backward		0.005***		0.006***		0.001
		(4.267)		(4.543)		(0.553)
Observations	2,336,260	2,336,260	1,625,472	1,625,472	708,417	708,417
R-squared	0.937	0.937	0.939	0.939	0.932	0.932
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Class by year FE	Yes	Yes	Yes	Yes	Yes	Yes
Clustered SE	Year	Year	Year	Year	Year	Year

Table B2: AI-Absorbing Innovation and Patent Value (alternate fixed effects)

Panel A presents results from estimating the following patent-level equation:

$$ln(\xi_p) = \alpha + \beta\ PctAIcite_p + \gamma_1 X_i + \gamma_2 \ln Backward_p + \eta_{it} + \omega_k + \epsilon_p$$

and Panel B presents results from estimating the following patent-level equation:

$$ln(\xi_p) = \alpha + \beta\ I_{AIcite_p} + \gamma_1 X_i + \gamma_2 \ln Backward_p + \eta_{it} + \omega_k + \epsilon_p$$

where ξ is the Kogan et al (2017) value of patent p, *PctAIcite* is the share of a patent's backward citations that cite an AI patent, I_{AIcite_p} is an indicator variable that takes the value one if at least one backward citation is made to an AI patent. *ln Backward* is the log of the number patents cited by the patent, X is a vector of firm-level control variables including log market capitalization and log idiosyncratic volatility, η_{it} are firm times year fixed effects, ω_k are technology class fixed effects. Standard errors are clustered by year. ***, **,* indicate significance at the 1%, 5%, and 10% respectively. All variable definitions are provided in Table XX in the appendix. The sample period is 1979 to 2024.

	(1)	(2)	(3)	(4)	(5)	(6)
	PANEL A: Dependent Variable in all columns: $ln(\xi)$					
	All patents		Non-IT patents		IT patents	
PctAIcite	0.000	0.000	0.010*	0.010*	-0.001	-0.000
	(0.068)	(0.141)	(1.949)	(1.964)	(-0.253)	(-0.163)
Market cap (logs)	0.885***	0.885***	0.890***	0.890***	0.867***	0.867***
	(35.229)	(35.227)	(35.138)	(35.137)	(28.325)	(28.321)
Backward		0.002**		0.001*		0.002***
		(2.610)		(1.844)		(3.416)
Observations	2,326,860	2,326,860	1,617,126	1,617,126	704,349	704,349
R-squared	0.964	0.964	0.964	0.964	0.962	0.962
	PANEL B: Dependent Variable in all columns: $ln(\xi)$					
	All patents		Non-IT patents		IT patents	
I_{AIcite}	0.003***	0.002**	0.006***	0.005***	0.003**	0.001
	(3.532)	(2.177)	(3.420)	(2.762)	(2.375)	(0.772)
Market cap (logs)	0.885***	0.885***	0.890***	0.890***	0.867***	0.867***
	(35.229)	(35.227)	(35.137)	(35.136)	(28.326)	(28.321)
Backward		0.001**		0.001		0.002***
		(2.121)		(1.495)		(3.116)
Observations	2,326,860	2,326,860	1,617,126	1,617,126	704,349	704,349
R-squared	0.964	0.964	0.964	0.964	0.962	0.962
Firm by year FE	Yes	Yes	Yes	Yes	Yes	Yes
Class FE	Yes	Yes	Yes	Yes	Yes	Yes
Clustered SE	Year	Year	Year	Year	Year	Year

Appendix C: Classification of Automation Patents and Process Patents

C.1. Classification of automation patents

We machine-read patent descriptions and classify a patent as an 'automation' patent if it meets any one of the following four criteria:

1. One or more appearances of the following words:
   - automatic or automatically
   - automate, automation, or automatization
   - mechanize or mechanization
   - robot, robotic, robotize, or robotization
   - labor-saving
2. To identify computer-assisted manufacturing: One or more appearances of the following bulleted list of words **and** the word *machine*, *manufacturing*, *equipment*, or *apparatus*
   - computer aided or computer-aided
   - computer assisted or computer-assisted
   - computer supported or computer-supported
3. To identify computer numerical control (cnc) patents which indicate automated control of machine tools: One or more appearances of the following words
   - numerically controlled or numerically-controlled
   - numeric control or numeric-control
   - cnc
4. To identify self-operated equipment: One or more appearances of the word 'self-' **and** the word *machine*, *manufacturing*, *equipment*, or *apparatus*

C.2. Classification of process claims

We follow Bena and Simintzi (2022, 2025) and take advantage of the standardized vocabulary with stilted legalistic terms used for process patents. Process patent claims often contain words such as 'method of', 'process for', to describe the steps or procedures involved in the invention. We classify a claim as process-oriented when its opening phrase matches the following pattern: an optional article ("a," "an," or "the"), followed by up to five optional descriptive modifiers, and then the keywords *method* or *process*. The expression also allows for hyphenated modifiers and flexible punctuation or spacing immediately following the keyword. Formally, this flexibly captures claim openings such as

- "a method…"
- "an improved manufacturing process…"
- "the computer-implemented method…," etc.

The matching pattern excludes occurrences in which *method* or *process* appears later in the sentence rather than as the claim's principal grammatical subject. Anchoring the pattern to the start of the claim text ensures that the classifier targets the conventional legal structure of process claims rather than incidental keyword usage elsewhere in the document.